Details of the L3 Wavefront Sensor for GLAS/NAOMI

Simon Tulloch
March 2006

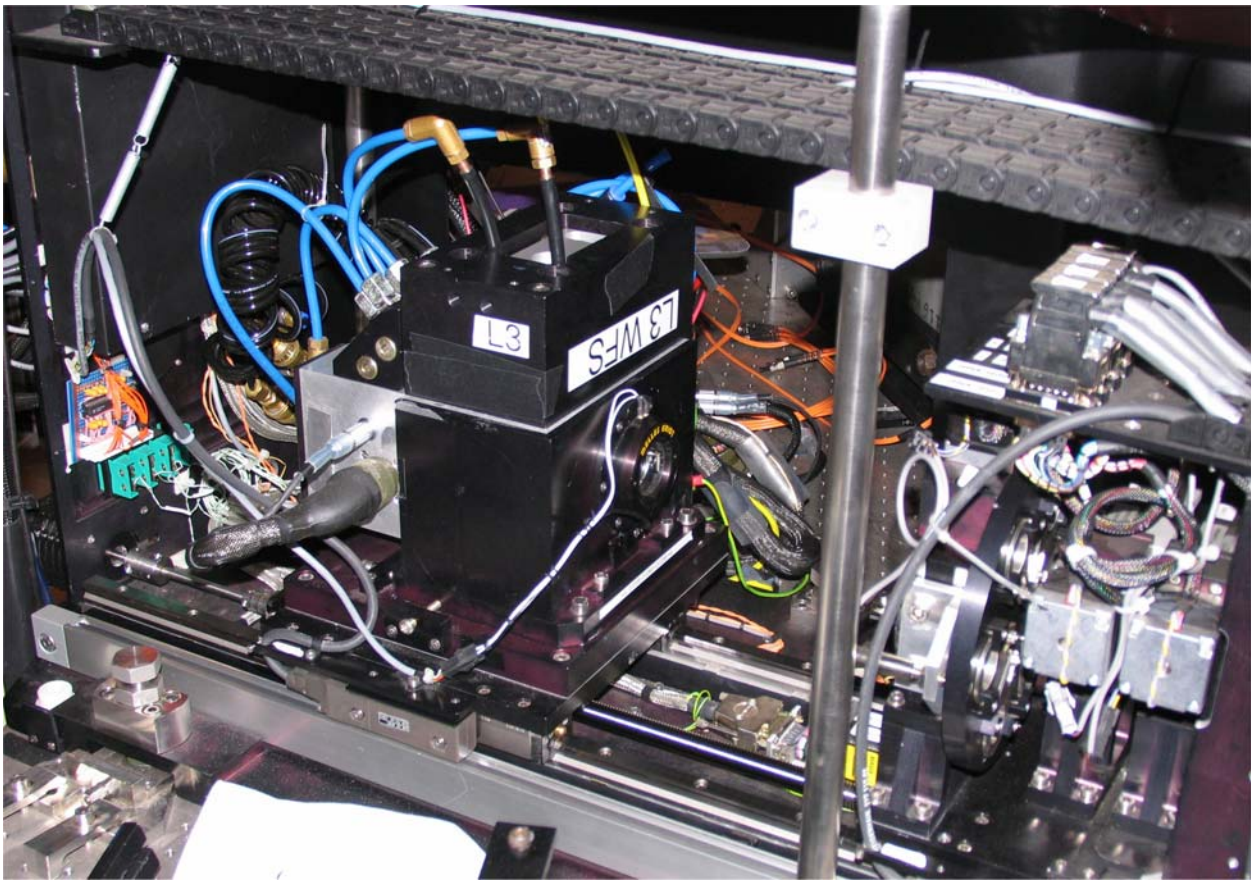

L3 WFS head mounted on NAOMI

1. Summary

This document is both a user guide and a technical reference for the electron multiplying (CCD60 L3) Wavefront Sensor (WFS) that has been fitted to NAOMI. This detector will be used for natural guide star wavefront sensing. It is primarily intended to provide tip-tilt (TT) correction for the GLAS Rayleigh AO system. In this application it is required to image a 17th magnitude guide star through a simple doublet lens. The sensor is also, however, applicable to Shack-Hartman (SH) sensing of a natural guide star through a lenslet array. It was first used on-sky in this mode in June 2005.

When used in TT mode with GLAS the camera is read out through a Linux PCi system. With a change of firmware in the camera controller it is also possible to read the camera out through the existing NAOMI VME system. Here the camera internally reformats its image so that it appears to the VME system as if it were the original CCD39 WFS.

In both applications it should give a considerable improvement in guide star limit over a conventional CCD.

2. Performance of the CCD60 L3 WFS Head.

2.1. Detector

Electron multiplying CCD60 with 128(+8 overscan) x 130 x 24μm pixels. Plate scale 0.25”/pixel. Integral Peltier cooler. Supplied by E2V Technologies. Operated using a GenII SDSU controller.

2.2. Readout Speeds

These are shown in the table below. When synched to the 5KHz Pockels cell the frame rate falls considerably.

Mode	Size (Pixels)	Frame time (un-synched)	Frame time (synched)
1	88 x 80	30.8 Hz	25.3 Hz
4	8 x 8(+2)	340 Hz	214 Hz
5	16 x 16(+2)	200.4 Hz	146 Hz
6	32 x 32(+2)	107.2 Hz	95.3 Hz

2.3. Dark Current

When the Peltier is running, dark current drops to 0.4e per pixel per second. This corresponds to a CCD temperature of about -28C. This was measured using exposures of 10s. The CCD60 is an AIMO device and has intrinsically low dark current. When running at 30 fps this will generate about 1e per row in Mode 1. When running at 340Hz in Mode 4 it will produce 1e in every 12 frames, a completely insignificant level.

2.4. System Sensitivity

When the L3 gain is unity, the system gain is 8.7e/ADU. As the L3 gain is turned up, this number will gradually fall. Without L3 gain the camera may give the impression of being rather blind. It is really designed to work at maximum gain (gain index=255) where the sensitivity is about 0.07e/ADU.

2.5. Linearity

The Image area full well is about 260Ke-. However, when L3 gain is applied, the output amplifier will begin to go non-linear long before the image area starts to saturate. The linearity seems to be fairly independent of the L3 gain and is good up to 25KADU. When attached to the NAOMI VME, pixel data is truncated to 14 bits.

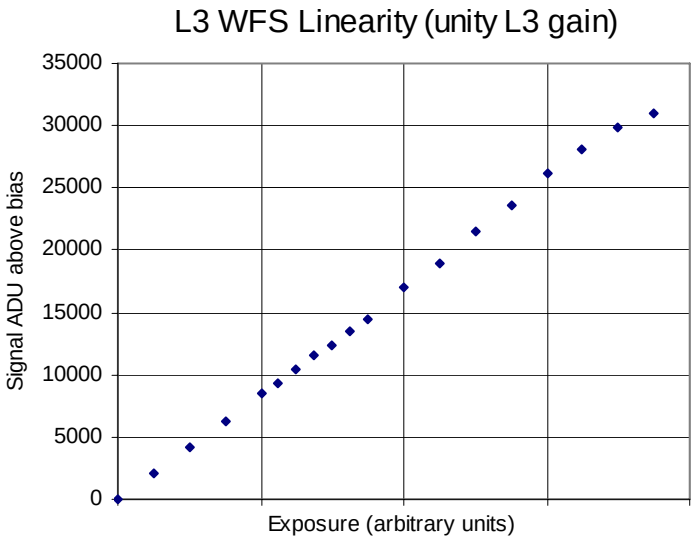

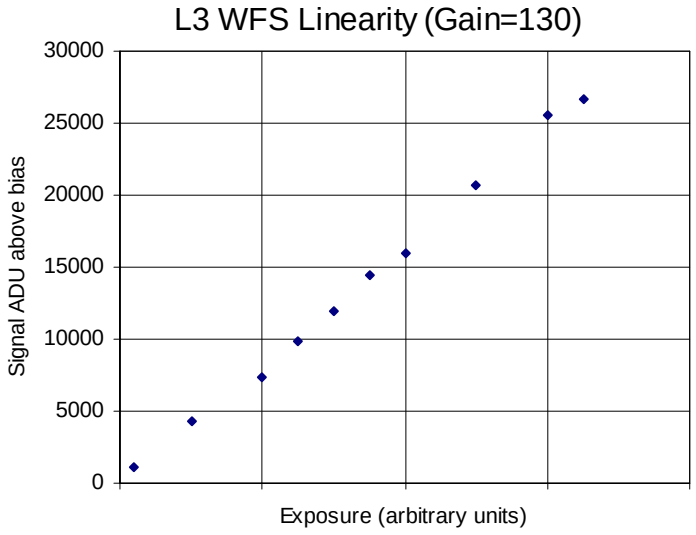

2.6. Read Noise

The inherent readnoise of the CCD is 54e RMS. Sub-electron noise therefore requires an L3 gain of greater than 54. L3 gains up to 212 have been tested. Higher gains should be possible if required but this will not reduce the read noise much further and will have the negative side-effect of reducing the dynamic range. With gain set to 212, the signal becomes non-linear at about 150 photons per pixel per frame. The high gain frames, one of which is shown below, show clearly individual photo-electron events thus demonstrating the sub-electron read noise.

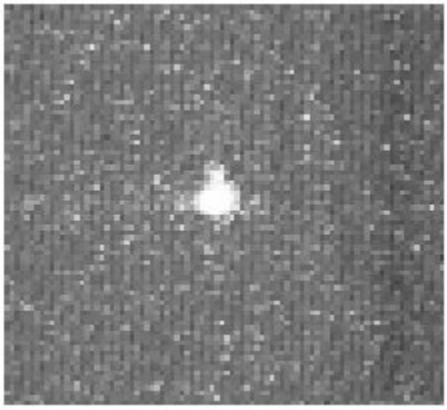

The cumulative noise in on-sky images was 0.74e RMS. This value is a combination of CIC, dark current and noise in the output amplifier.

2.7. Quantum Efficiency.

The following data is from E2V data sheet. Their QE data is normally very reliable.

400nm	500nm	550nm	600nm	650nm	700nm	800nm	900nm	950nm
50%	90%	92%	92%	91%	87%	70%	42%	27%

2.8. Frame Transfer Time

It takes 170us to move the image area pixels under the shielded store area. This is about 7% of the minimum exposure time. Some vertical smearing of very bright guide stars may therefore be visible.

2.9. Optimal Focus

The conventional CCD39 WFS has optimum focus at 211 units (on the NAOMI engineering GUI). The L3 WFS has optimal focus at 170 units.

2.10. Charge Transfer Efficiency

This appeared to be good in both axes. Vertical transfer appears perfect but the Horizontal transfer shows a slight trail at high gains as is normal with L3 CCDs. A spot image consisting of 5 summed frames shows this effect below :

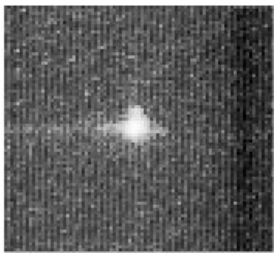

2.11. Clock Induced Charge (CIC)

This noise source is generated by fast clock transitions in the CCD. It is present in all detectors but rarely noticed. In L3 sensors where individual electrons become clearly visible, it becomes a problem. To measure this noise source it is necessary to combine many bias frames, at least 50. Each frame is individually histogrammed to determine its true bias level. This is taken as the peak of the histogram. Each frame then has its bias level subtracted and the resultant images averaged. This mean signal in the resultant image will then be a combination of dark current and CIC. Clearly a small error in the bias determination can produce a huge error in the CIC measurement so it is important that it is done carefully. A C program was written to do this using a two step iterative process. On the first pass each input image was coarsely histogrammed to locate the approximate position of the peak (corresponding to the bias level). In the second pass the program 'zoomed in' on this peak to resolve it to an accuracy of 1 ADU.

It is then necessary to convert this result into electrons per pixel per frame by dividing by the system gain(ADU/e^-). This is best done using a new technique developed at ING that relies on examining the event height distribution of the photoelectron or CIC events. A large stack of images is needed to give good statistics. The images need to contain a very low level of illumination or CIC that has been generated in the Image Area (which are indistinguishable from real photo-electrons). The pixels in the stack of input bias subtracted images are then histogrammed and graphed with a LOGe vertical scale. The system gain in ADU/e^- is then simply the negative reciprocal of the graph's gradient. For this camera a good value for the gain was arrived at by just using bias frames containing CIC. It was important when measuring the histogram slope to ignore values close to the bias since the Gaussian read noise disturbs the statistics in this region. This can be seen in the following graph where there is a slight 'hump' at lower values due to the amplifier read noise. The result of 14.2 ADU/e^- was in good agreement with photometric tests that compared the system gain of the two WFS heads whilst observing an illuminated pinhole.

A C program to calculate CIC and system gain is available on the 'develectronics1' Linux PC at `/home/smt/WFS/calcL3noise`. The program takes a text file containing a list of input Mode1 fits images as an input parameter. It outputs a bias subtracted image sum and a histogram of the image area pixels.

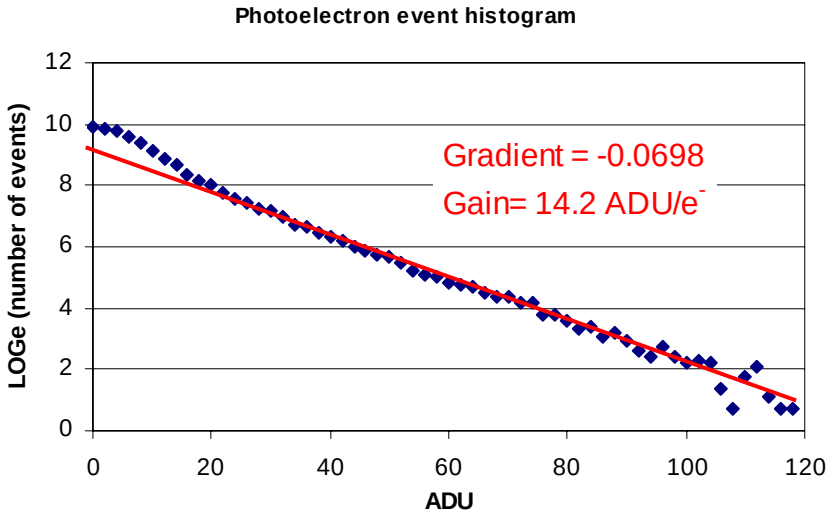

Using this result the CIC+dark current was then determined to be 0.3e^- per pixel per frame. The bias frames had an effective integration time of 30ms so the dark current contribution to this sum would have only been 0.01e per pixel per frame. This level of CIC is acceptable but could probably be improved upon by careful attention to the shape of the clock edges. Reduction of clock amplitude can also reduce CIC but in this case the clocks were already at the minimum amplitude compatible with good charge transfer efficiency.

2.12. Comparison with the conventional CCD39 WFS

Firstly it was necessary to check the optical throughput of the two WFS arms. This was done by observing a fairly bright pinhole through both cameras with the same exposure time. The total detected flux in the two pinholes was measured using IRAF and converted to electrons. The conversion factor of the CCD39 was 0.56e/ADU, that of the CCD60 was 0.07e/ADU. The result was that with the pinhole iris set to -300 units the CCD60 saw 29151 electrons per frame, the CCD39 saw 32430 electrons per frame. The difference can be explained by small errors in the gain measurements and the QEs of the two chips. This demonstrated that the two detectors have throughputs within 10% of each other. The pinhole was then further stopped down and comparison images taken with both cameras. Two such comparison images are shown below and it is clear that the L3 camera is more sensitive.

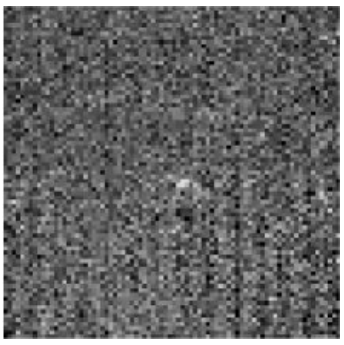

Faint pinhole with CCD39 WFS

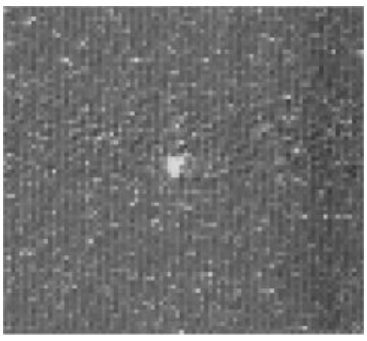

The same pinhole with the L3 WFS

This improved sensitivity is demonstrated below more quantitively. The graphs show cuts through the peaks of the pinhole images. This was done at two separate iris values to show the how as the conventional CCD image sinks into the noise the L3 CCD still produces high SNR data. The vertical scale has been converted to electrons. The CCD39 shows a higher peak because it was better focused. When the total flux in the two images is summed and compared they were found to be the same, so the problem was purely one of focus rather than sensitivity.

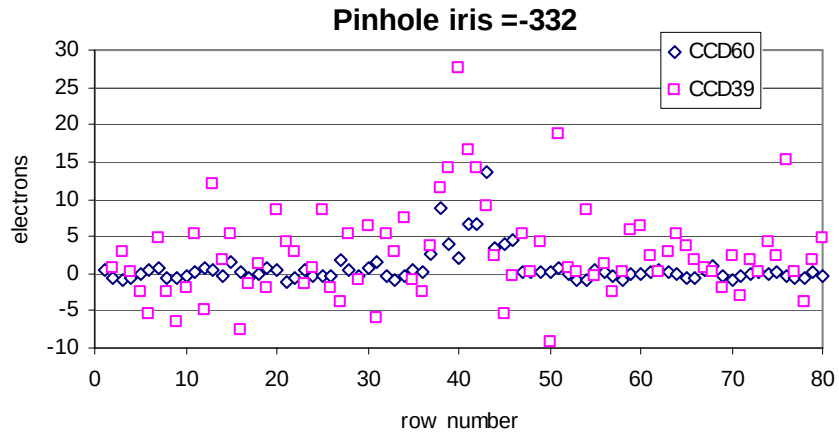

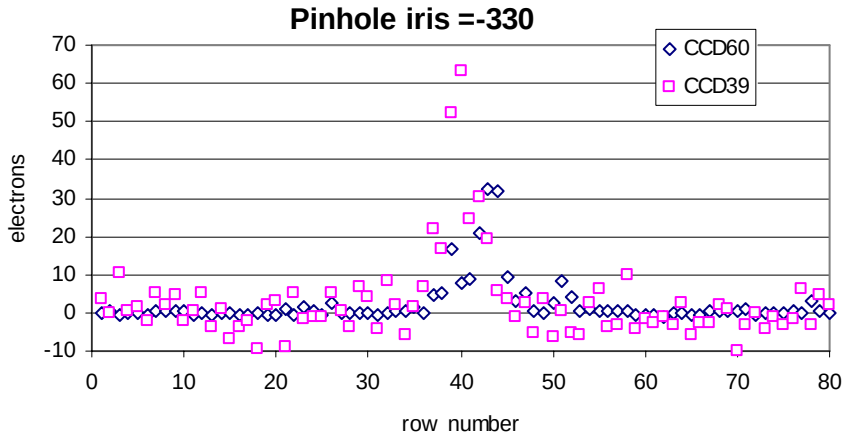

With the pinhole set to -332, the image clearly rises above the noise on the L3 camera but is barely detected on the conventional camera.

The L3 camera was also calibrated using a star of known brightness. A M_R =14.7 star at an altitude of about 50 degrees produced a signal of 36Ke$^-$ per second. This measurement was done with no filter but with the 50:50 dichroic in position. Astronomy Group models predicted 42 Ke$^-$ per second so this is a good confirmation of system throughput.

3. Engineering Details.

3.1. Light Tightness of the Enclosure.

With the beam-splitter cube in the out position and the Grace lights on the light leakage into the L3 WFS amounts to 2.6K photo-electrons/s. With the beam-splitter in, this rises to between 8K and 20K photo-electrons/s depending on pixel position. Even with the beamsplitter out, the GRACE lights are sufficient to trigger the L3 overexposure protection circuitry (more on this later).

3.2. Image Subtleties

When pixels are skipped rather than read, the L3 multiplication gain is slightly lowered. This means that small windows show a slight reduction in gain (20%) compared to full frame images. This is thought to be due to the mark space ratio of the multiplication clock being higher during skipping than pixel reads. This causes an increase in current supplied by the clock board and a slight reduction in voltage output. Only a very small reduction in voltage can cause a large change in L3 gain.

Initially, quite serious bias gradients were seen in the y-axis. Further investigations suggested that these were due to the L3 gain at the start of a frame being much higher than normal. The CIC in the image was thus preferentially amplified in the first few rows giving the impression of a raised bias. The effect was to severely offset the guide star centroids and prevent detection of faint stars. The cause was the same as for the gain reduction during pixel skips : a change in duty cycle of the multiplication clock. Immediately prior to the controller reading the image window one has to consider that the chip has spent some time integrating and then doing a frame transfer of the image. During this time the multiplication clock was idle i.e. supplying no current. At the start of pixel readout the clock will then start up with a slightly elevated value which will settle down after a few rows to a more stable lower value. The solution to this is to make sure that the HV clock is always busy, even during integrations and vertical readouts. Some experimentation was required to get the tuning right. In addition, the L3 register was fully flushed using the pixel skip waveform immediately after the vertical skip to the readout window had completed. These two modifications completely removed the vertical bias gradient.

A further bias gradient of lower amplitude was seen in the horizontal axis and was again due to cadence changes in the HV clock pulses. It was removed by reading additional warmup pixels before the start of the image window within each row.

3.3. Beam Switching between L3 and conventional WFS heads

There was originally a half silvered mirror to allow light to be shared equally between master and slave WFS heads. Since the Slave was removed and replaced with the L3 head, this half silvered mirror was replaced with a fully silvered equivalent. It is on a manually operated slide, which rather confusingly has 4 detented positions. These are indicated in the diagram below:

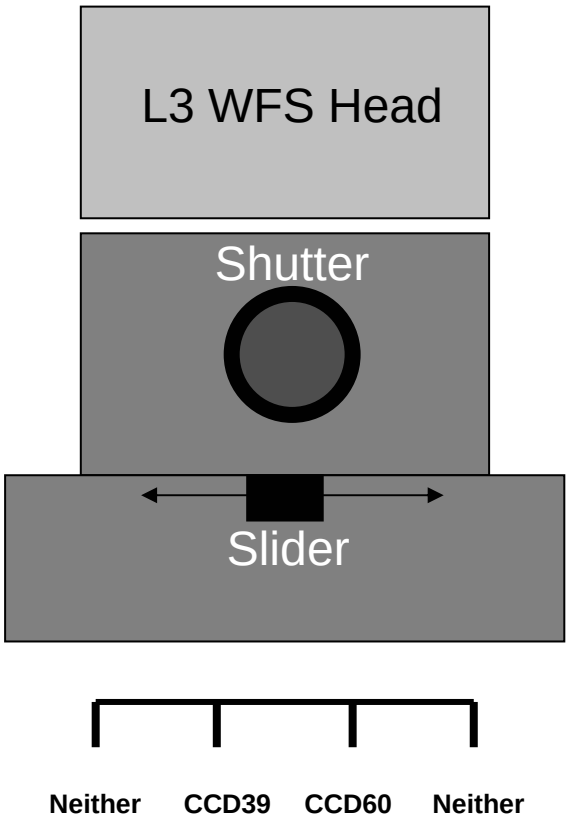

3.4. L3 Multiplication Gain Changes

The original Naomi system allows two readout speeds to be selected for the WFS from the TopGui ; Fast and Slow. The L3 camera reads out at the same rate whichever of these two is selected. Instead, the speed change is used to change the level of L3 gain. Slow speed gives a gain of 1, fast speed a gain of up to several hundred. The low level commands used to change speed are SLW (slow) and HIH (fast). These must be followed by the SYC command to effect a change. These commands are implemented on both VME and PCi versions of the L3 WFS code.

In addition, the PCi code contains an extra command ‘L3G’ that is followed by a gain level parameter that allows much finer setting of the L3 gain. The parameter is a number between 0 and 255, for example :

‘L3G 0 ‘ has the effect of setting the gain = slow speed gain
‘L3G 255’ has the effect of setting the multiplication gain = fast speed gain

Use of intermediate parameter values sets the gain to values between these extremes, however, there is not a linear relation between the two since the value actually sets the voltage of the multiplication clock rather than an exact gain level. L3G takes immediate effect and does not require a subsequent SYC command. It is hoped that this command can be used as part of an automatic gain control system.

The graph below shows the approximate variation in L3 gain with the value of the gain index. Gains higher than about 212 were not used, although gains of several thousand have been demonstrated with other L3 CCDs at ING.

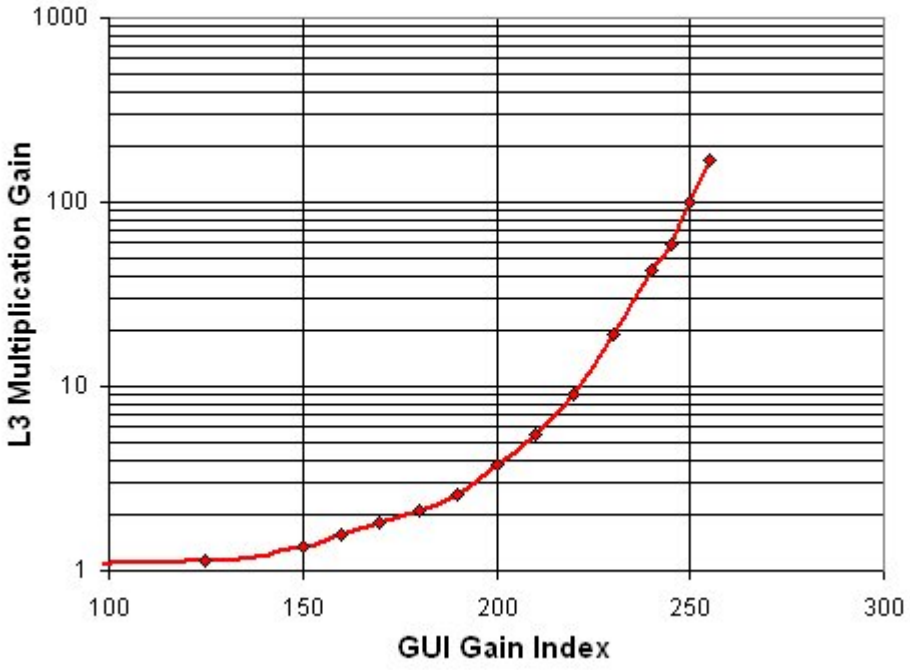

3.5. CCD Spares

We have a total of 4 CCD60 devices purchased or obtained free for the project. In addition to the operational device mounted in NAOMI, there is a second Peltier packaged CCD60 which is a drop in spare should the original fail. There is also an engineering grade ceramic packaged CCD60 for lab work where cooling is not required and an additional science grade ceramic packaged CCD60 with no immediate application.

3.6. Optical Orientation of the L3 WFS Head.

The CCD60 in the L3 WFS head was oriented within its head so that its columns lie in the same plane as the columns in the CCD39 Master Head. Any translations of a star image up and down a column on the CCD39 will also cause an up-down translation on the CCD60, although not necessarily in the same direction. When viewing images on TopGui, the CCD39 and CCD60 images do have the same orientation. The VME based Tip Tilt modes introduce a flip in both X and Y directions to the CCD60 images in order to get this correspondence.

When using the PCi TipTilt engineering GUI, the displayed images have the orientation shown below. If the images are saved as fits files from this GUI, they are displayed with a vertical flip.

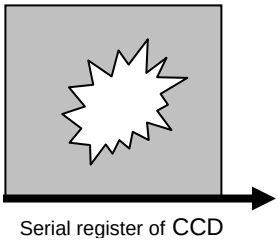

TipTilt GUI displayed image orientation

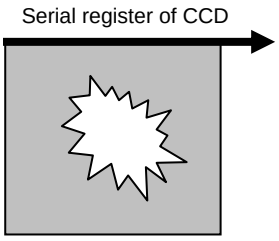

TipTilt GUI saved fits file orientation

The CCD60 silicon is 3.1mm (optically not mechanically) behind the front of the package cover glass. The corresponding value for the original CCD39 Master head is 3.5mm. The image below show the CCD package orientation within the L3 WFS head.

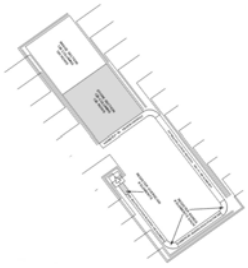

The diagram at right shows the serial register orientation on the L3 CCD, reoriented so as to correspond with the adjacent photograph.

3.7. L3 CCD Pipeline Details

The L3 Camera contains two end-to-end pipelines that greatly complicate the readout code.

L3 Pipeline : equal to 4 image rows. This is dealt with by keeping a count of rows read and only activating pixel digitisation when the pipeline is full. The pipeline is emptied at end of *each* frame by reading an extra 4 rows, so the readout speed does suffer.

ADC pipeline: equal to three pixels internal to the ADC chip. The only way to deal with this without causing image artifacts is to read out an extra three pixels at end of *each* row. The first three pixels at start of the subsequent row are disregarded as junk. Again, readout speed suffers but not greatly.

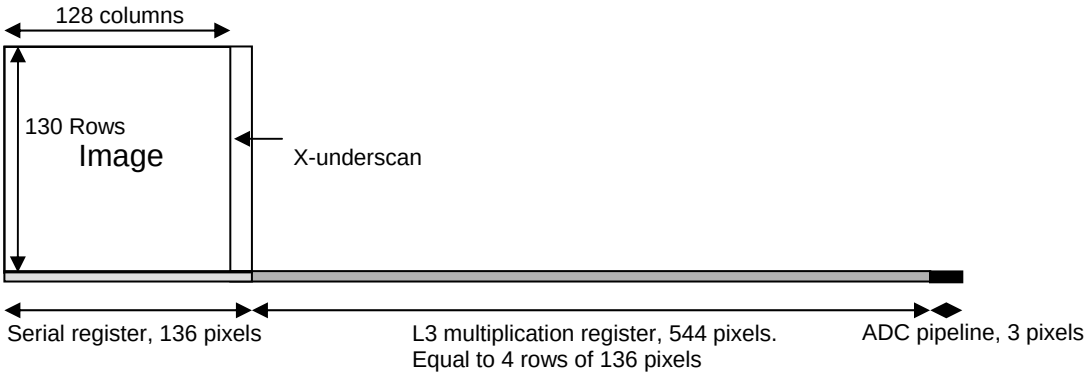

3.8. Use as a Shack Hartman Sensor

If readout mode 1 is selected the camera can read out a full image of the SH lenslet array. The frame rate is 30.8Hz. An actual Mode 1 image is shown below.

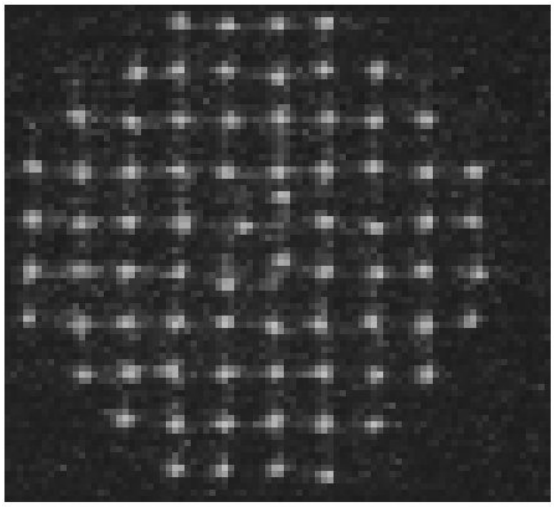

3.9. Changing the Window Position of the L3 WFS

If at anytime it is necessary to remove the L3 WFS head, it is unlikely to be remounted to a pixel level precision. There are adjustment screws in the head but it is far easier to simply change the position of the software defined window on the CCD. The fact that the CCD60 is bigger than really needed then becomes very useful. The following procedure should be used if the centre of the window needs to be realigned.

First select mode 1 readout for the WFS to give the largest field of view. Observe the NCU pinhole through the doublet lens to produce a spot on the WFS. Measure the spot image position on the TipTilt GUI and calculate the X and Y offsets from the centre pixel. The image displayed on this GUI is displayed with the readout direction to the right. The vertical sense is correct with the serial register at bottom. If the spot has a +ve horizontal offset then the value of XWINCENT will need to be *reduced* by that amount. If the spot has a +ve vertical offset then YWINCENT will need to be *increased* by that amount. These constants are defined in:

/home/dspdev/L3NAOMIPCi/ap_boot/timing_naomi_L3_header.asm

Once this file has been edited, the boot and application code will then need to be reassembled and burned into the Timing card EEPROM. This assembly can be done with a single command : '/home/dspdev/L3NAOMIPCi/assall'. Details of how to burn the EEPROM are provided elsewhere in this document.

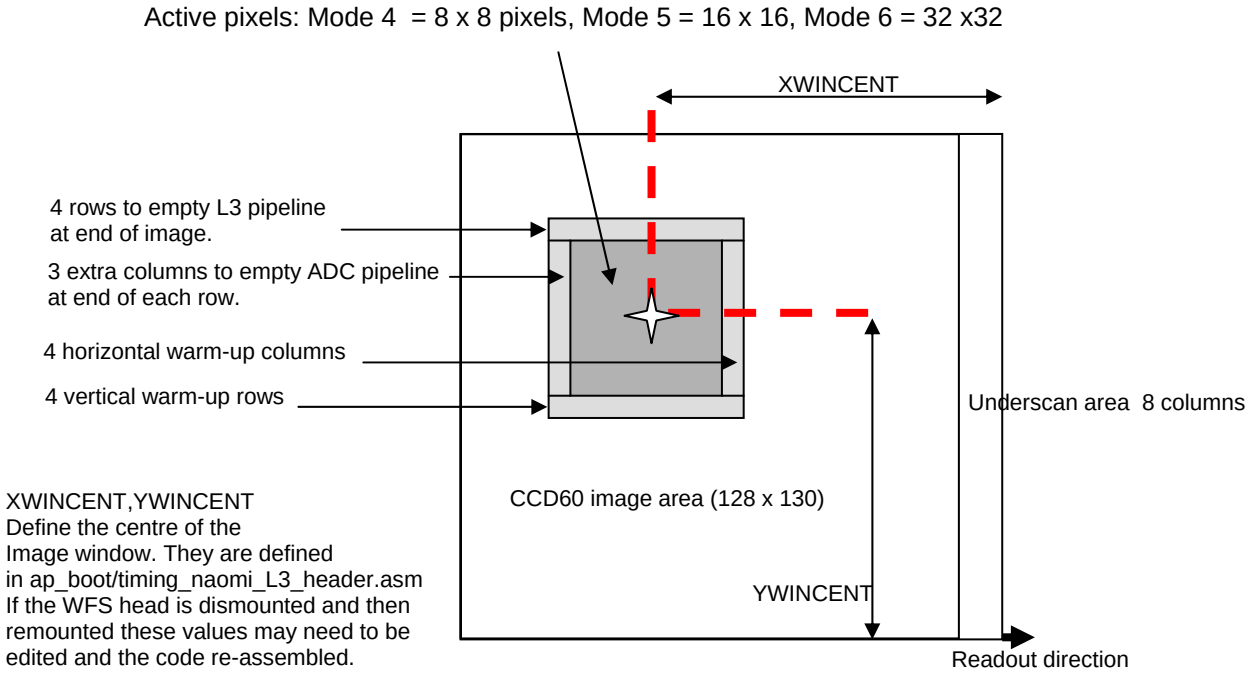

XWINCENT is nominally 61 , YWINCENT is nominally 87

3.10. Details of CCD60 Pixel Readout

Reading small windows on CCDs almost always leads to spuriously bright columns and rows at the extremes of the image. Bias gradients are also a common problem. These totally confuse the centroiding algorithms used to calculate the tip-tilt errors and it is essential that the CCD60 gives clean, flat images. The solution to this was to read extra columns and rows prior to an image window to 'warm up' the system. These extra columns were then edited out and discarded prior to the image being transmitted from the SDSU controller. This all required the creation of an input and output image buffer within the controller. As the development progressed it became apparent that buffering the images in this manner had many useful benefits. It allowed the transmitted image to have any format, the CCD60 could even be made to look as if it were a CCD39 and so retain compatibility with the NAOMI VME system.

For the Tip Tilt PCi system non-image pixels were superimposed upon the main image to give extra information. One pixel in each corner contained the electronic bias level of the chip. These 'bias pixels' were obtained by averaging one row of pixels that had been read from a previously flushed L3 register. They therefore contained no photo-charge from the Image Area and represented a true bias level. Additionally, the first and last rows of the image were taken from rows that actually lay 10 and 11 rows above the top of the image window and therefore represented a true background signal. These pixels, although containing clock induced charge, and dark current, would not be contaminated by any light from the guide star due to their remoteness from its image. These 'background rows' were needed for proper functioning of the centroiding algorithm. In mode 1, the non-science mode, it was not necessary to modify the first and last rows since they were already far from the guide star on account of the large dimensions of the image.

The images below shows the various regions labeled.

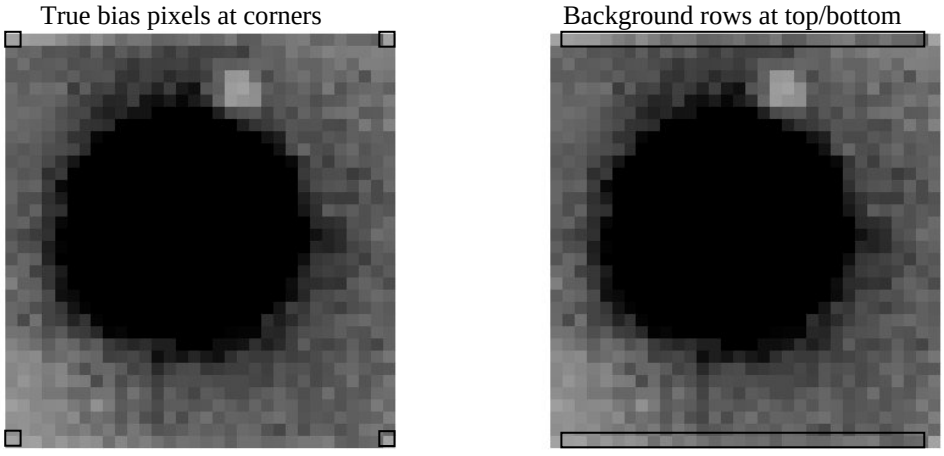

The diagram on the following page shows how this reformatting is achieved inside the controller.

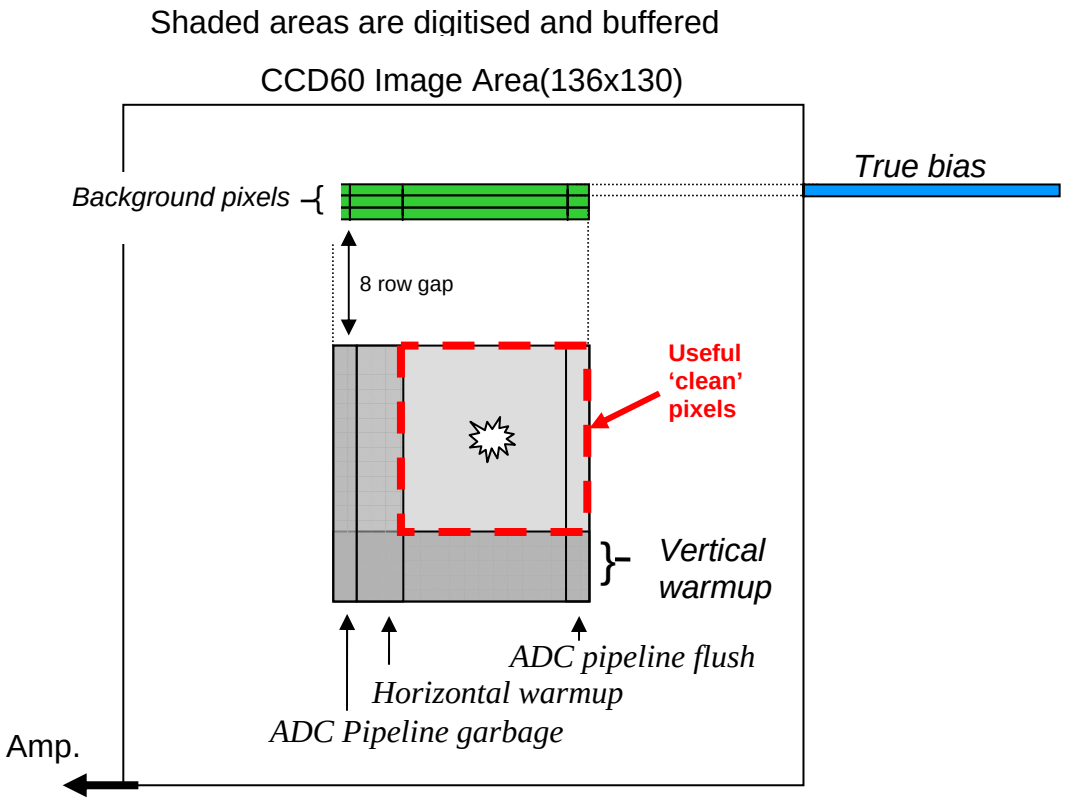
Shaded areas are digitised and buffered
CCD60 Image Area(136x130)
True bias
Background pixels
8 row gap
Useful 'clean' pixels
Vertical warmup
ADC pipeline flush
Horizontal warmup
ADC Pipeline garbage
Amp.

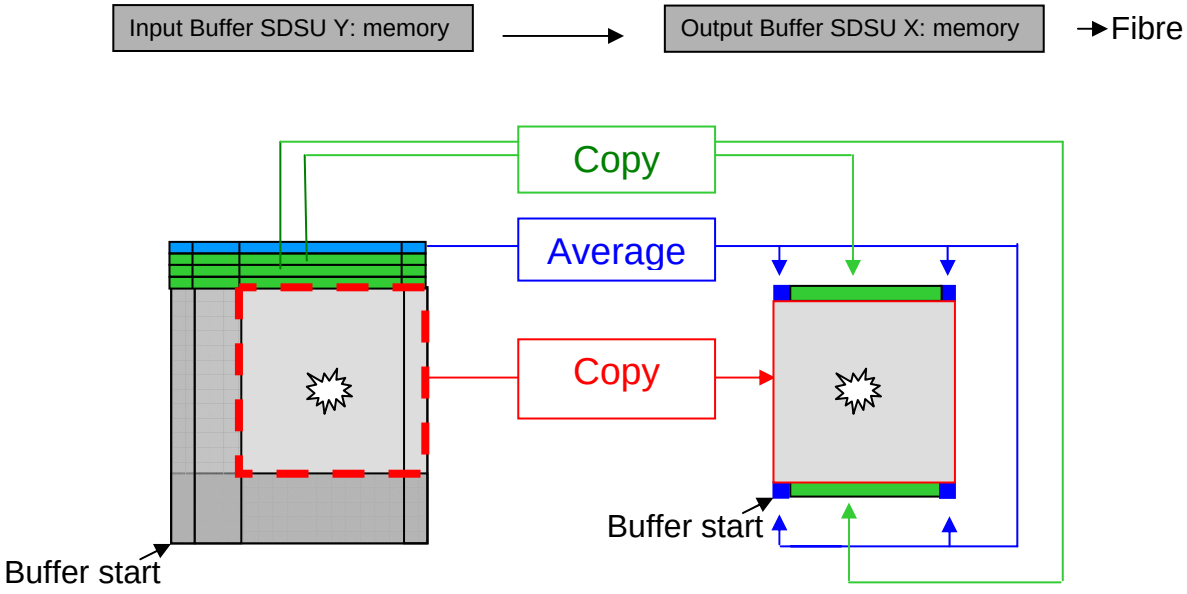
Input Buffer SDSU Y: memory
Output Buffer SDSU X: memory
Fibre
Copy
Average
Copy
Buffer start
Buffer start

3.11. CCD60 Peltier Package

The CCD came with its own hermetically sealed peltier package. Data supplied by the manufacturer recommended operation at 1.5Volts/5.5 Amps to obtain maximum cooling of about 45 degrees centigrade. A large linear supply was purchased for this purpose; switched supplies had previously been found to give a great deal of pickup noise when used with the ING autoguider peltier coolers at high current. The voltage output at the supply was 1.44 V which gave a Peltier current of 3 Amps and this was found to cool the chip to -28 degrees C in the GRACE laboratory environment. The warm side of the Peltier was cooled with the same loop as the conventional CCD39 head. There is, however, no interlock to cut Peltier power if the coolant flow stops. Since the dissipation is only about 4.5W, an interlock is not really necessary.

The graph below was supplied by E2V. It is from a theoretical model.

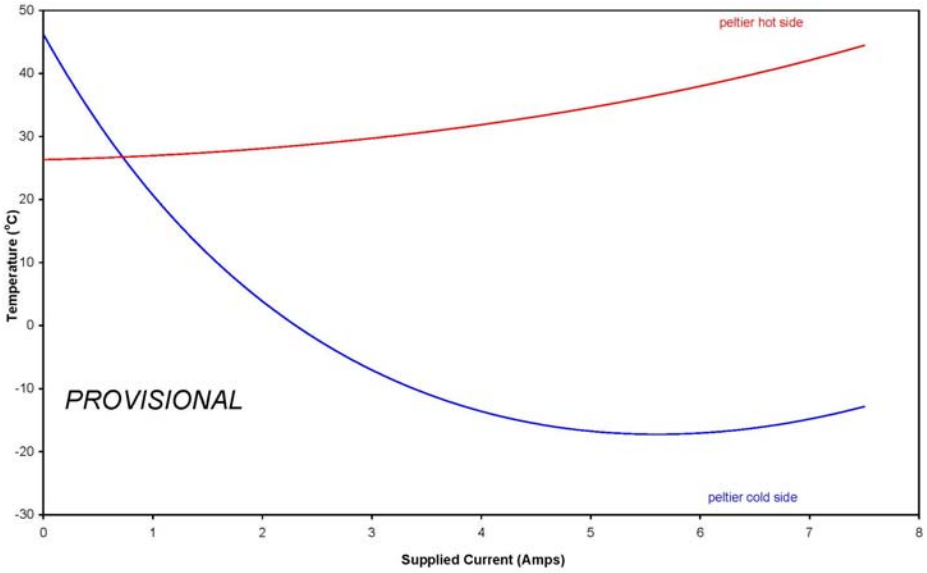

4. SDSU Camera Controller Software

The code for the controller was entirely resident in EEPROM. It was not necessary to download code (in the form of .lod files) over the fibre. Two EEPROMs are provided, one for connection to a PCi system, the other for connection to the original NAOMI VME.

4.1. PCi SDSU L3 Code

The main difference between the VME and PCi versions of the code is the need to send the RDA and IIA commands. IIA is transmitted by the SDSU controller to the PCi interface to zero the pixel counter. This is done before the very first frame after a mode change and every 32 frames thereafter. A delay of 100us is required after IIA is sent for the PCi to do its preparations. The RDA command must be sent by the SDSU controller to the PCi interface before each frame is transmitted. It has two parameters whose product is equal to the pixel count of each image. This prepares the PCi board to receive pixels. Included in this count are 10 extra header 'pixels' , 2 'footer' pixels since these words (that contain frame diagnostics rather than pixel values) are transmitted with bit WW of X:PBD set. This configures the fibre interface to transmit 16 bit words and means the PCi interface thinks they are true pixels. With bit WW of X:PBD clear the interface uses 24 bit words and expects commands rather than pixels to be transmitted and received. The minimum time between RDA being sent and pixels being sent is 600us, the maximum delay is approximately 3ms after which a timeout occurs. Simply adding a delay loop of say 650us after the RDA command is not acceptable since it will greatly slow readout, instead the CCD readout itself (or parts of it) are used to provide the delay, remember that pixels are all buffered internally before being transmitted.

In the 8x10 TT mode the RDA command is sent when the L3 pipeline is filled. From this point on the CCD readout provides the ideal delay time. If any analogue tuning of the CCD in the future causes the readout time to slow, care must be taken that this delay does not rise above 3ms, after which readout will certainly fail.

In the 16x18 TT mode the image is read in about 4ms, too long to allow the RDA to be sent at the start of readout. In this mode the only solution was to send RDA after image area readout but before the image was reformatted in the controller and the bias rows added. These latter two operations only occupy about 500us so, unfortunately, it was necessary to add an additional 160us delay loop with a direct effect on readout speed. One other solution may be to send RDA after say 50% of the image rows have been read, so allowing the extra delay loop to be removed, but this would affect the cadence of pixel readout with the likely introduction of image bias level artifacts.

In 32x34 TT mode a similar method is used to get the correct RDA delay. Since the image is larger and reformatting takes longer the extra delay loop that needs to be added is correspondingly smaller and there is only a small impact on readout time. SEND_IIA and SEND_RDA are defined in the slow bootcode.

4.2. VME SDSU L3 code

The code is modified from the original ATC CCD39 code. Almost all the bootcode is the same although of course the part that deals with the actual CCD readout is largely new. The ADC pipeline in the SDSU controller was not properly dealt with by the original ATC code. Their solution was to read out an extra 2 pixels from each corner of the chip at the end of each frame in order to flush the pipeline. The first 8 pixels at the start of each frame are therefore pipeline junk and are disregarded further downstream. The L3 VME code also transmits these 8 junk pixels at the start of each frame in order to retain compatibility. There appears to be some confusion over the actual size of the ADC pipeline since the CCD60 code would only give clean images if it was assumed to be 3 pixels long rather than the 2 pixels assumed by the ATC. Even though the L3 WFS is to be used on-sky with small windowed TT modes, the camera must also contain an implementation of Mode 1 if it is to work through the VME system. Mode1 is the full frame CCD39 mode that produces a 88 x 80 pixel image. It is defaulted to when the camera system is initialized. In this mode the CCD60 labors to give more than 30 frames per second which is probably too slow to be of use. Nevertheless it must be present or the system will not initialize. Mode 1 was difficult to write since the VME expects quadrant style data from a four quadrant CCD39 chip. 'Quadrantising' of the CCD60 format image into the format expected by the VME system was done internal to the SDSU controller and used up almost all the available memory as image buffers. Quadrantising was not required by the new TT modes (4,5,6) since the higher level code was programmed to expect plain format data.

4.3. L3 WFS Source Code Location/Structure.

All the code can be found in the `/home/dspdev` account in the following directories :

`/home/dspdev/L3NAOMIVME` : VME applications

`/home/dspdev/L3NAOMIPCi`: PCi applications.

`/home/dspdev/L3VOODOO`: Code that was used for fine tuning of the readout using the Voodoo program, which provides a useful test system and means for storing images. The files to be utilised are timing.lod files that are downloaded into the controller in the standard way. When using Voodoo a standard timing card EEPROM must be fitted in place of the WFS EEPROM. The corresponding utility.lod download file can be found in `/home/dspdev/util/voodoo`. This directory contains a particularly useful sub-directory called `FULLFRAME.` The code in this directory reads out the full CCD60 (136x150 pixels) which can be used for optical alignment of the head.

`/home/dspdev/L3NAOMI87`: Modified L3VOODOO code that allows a CCD87 engineering CCD to be attached in place of the CCD60 for lab based tests. An engineering grade CCD60 was later obtained so the code in this directory is no longer that useful.

Make files are not used. Instead the assembly/linking is done by an 'ass' script in each directory.

4.4. WFS SDSU code assembly

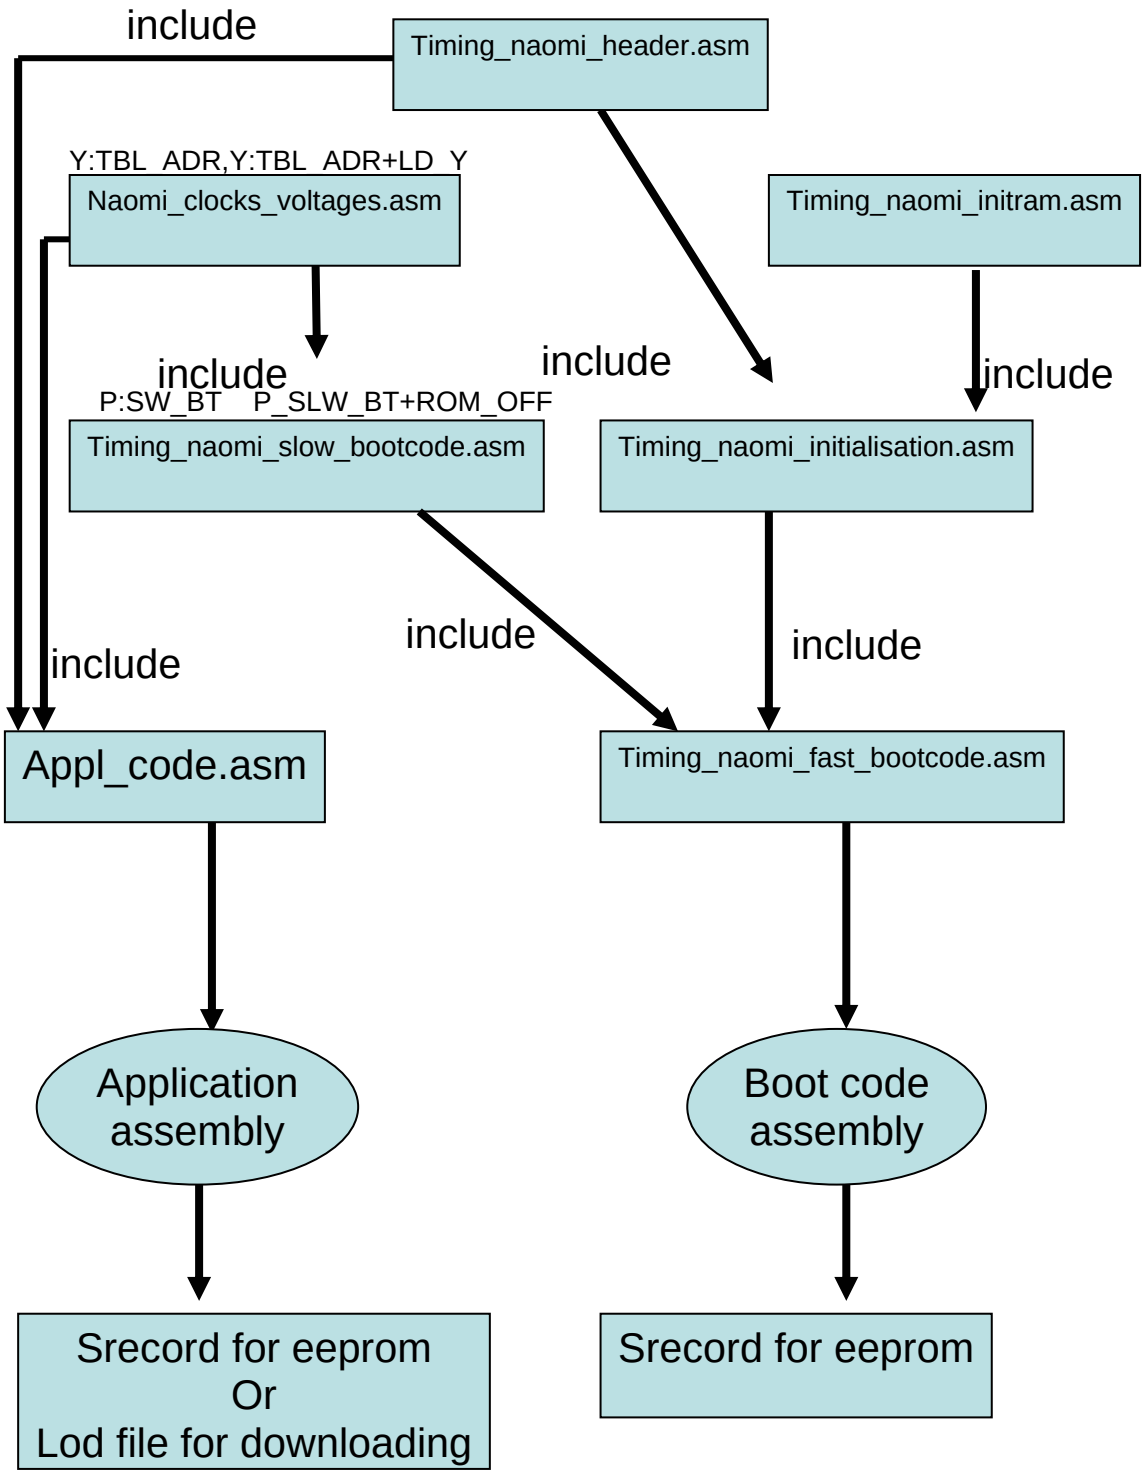

4.5. Attaining Maximum Readout Speed

The readout of the CCD60 is limited in speed by the SDSU controller. The chip itself can go at 1000fps full frame, equal to about 16MPix/s. The SDSU runs into a limit at around 1MPix per second due to its ADC conversion rate of 1us but also due to the speed at which clock waveforms can be sent to the clock card. There is now an option to retrofit a 10MHz 14 bit converter to the standard Video processor card. This converter, the ARC-47, is sold by Astronomical Research Cameras and ING already has one for purposes of experimentation. Adding one of these to the WFS system may not, however, make much difference since the conversion is done in parallel with CCD readout and about 1us per pixel is required by the video processor anyway , to get an adequately low noise level. The place where more significant time savings can be made is in the sending of clock waveforms to the clock driver board, in particular the clock sequences that govern the skipping of unwanted pixels to the left and right of the image window. It is these skips that occupy the greater part of the readout time. The standard way to send clock sequences is to use the CLOCK subroutine, passing the address of the relevant waveform table as a parameter in RO. The subroutine is thus called once for *each* pixel that is skipped and creates a lot of overhead. The DSP assembler code for this is shown below:

```
; STANDARD METHOD OF SKIPPING A PIXEL
;example code to skip 100 pixels
            DO    #100,SL1    ; Skip 100 pixels
            MOVE  #<SSKIP,R0  ; Address of serial skip waveform
            JSR   <CLOCK      ; Jump to subroutine
SL1         NOP


; Core subroutine for clocking out CCD charge
; R6 already contains address of the clock board
CLOCK       MOVE    Y:(R0)+,X0      ; # of waveform entries
            MOVE    Y:(R0)+,A       ; Start the pipeline
            DO      X0,CLK1                 ; Repeat X0 times
            MOVE    A,X:(R6) Y:(R0)+,A      ; Send out the waveform
CLK1        MOVE    A,X:(R6)        ; Flush out the pipeline
            RTS                     ; Return from subroutine

; Waveform table to skip a pixel, stored in Y memory
; First entry is the table length.
SSKIP       DC    SSKIP_END-SSKIP-2
            DC    VIDEO+$000000+%1110100
            DC    CLK2+S_DELAY+PhRH+H1LH+H1RH+H2LH+H2RH+H3L
            DC    CLK2+S_DELAY+PhRH+H1LL+H1RL+H2LH+H2RH+H3L
            DC    CLK2+S_DELAY+PhRH+H1LL+H1RL+H2LH+H2RH+H3H
            DC    CLK2+S_DELAY+PhRH+H1LL+H1RL+H2LL+H2RL+H3H
            DC    CLK2+S_DELAY+PhRH+H1LH+H1RH+H2LL+H2RL+H3H
            DC    CLK2+S_DELAY+PhRH+H1LH+H1RH+H2LL+H2RL+H3L
SSKIP_END
```

Avoiding the overhead from the subroutine jumps can improve readout speed by 30%. This can be done by ‘unwrapping’ the loops so that the waveform table is 100 times longer (in the above example) and CLOCK called only once. This is however very wasteful of memory and the SDSU controller has very little to spare, especially internal fast access memory. The better alternative is to keep the waveform table 1 pixel long but to access it using ‘Modulo

Addressing'. This is a hardware feature of the processor and allows the waveform table to be repeatably looped through using a pointer register whose value is incremented after each pixel.

The DSP assembler code for this is shown below:

```
; NEW FASTER 'MODULO' METHOD OF SKIPPING A PIXEL
; Skip 100 pixels
            MOVE  #(SKIPLENGTH*100-1),B0
            JSR   <SKIP_PIXELS
            NOP


;Subroutine to skip a block of pixels. Called only once per block
; i.e. twice per line of image, once to left and right of image window.
;----------
SKIP_PIXELS
;----------
; B0 parameter=(number to skip x length of skip routine) -1
            MOVE  #<SSKIP,R0              ; set up circular
            MOVE  #(SKIPLENGTH-1),M0      ; waveform table
            NOP                           ; M0 sets modulo mode
            MOVE  Y:(R0)+,A               ;
            REP   B0
            MOVE  A,X:(R6)  Y:(R0)+,A     ; Reads next table entry
            MOVE  A,X:(R6)                ; and  transmits  to clock board.
            MOVE  #$ffff,M0               ; Return to normal addressing
            RTS                           ; i.e. 'Non-Modulo'

; Waveform table in Y memory
SSKIP       DC    VIDEO+$000000+%1110100
            DC    CLK2+$10000+PhRL+H1H+CEH+H2L+H2HVH+H3L+RDCH
            DC    CLK2+$00000+PhRL+H1H+CEH+H2H+H2HVH+H3L+RDCH
            DC    CLK2+$00000+PhRL+H1L+CEL+H2H+H2HVL+H3L+RDCH
            DC    CLK2+$00000+PhRL+H1L+CEL+H2H+H2HVL+H3H+RDCH
            DC    CLK2+$00000+PhRH+H1L+CEL+H2L+H2HVL+H3H+RDCH
            DC    CLK2+$10000+PhRL+H1H+CEH+H2L+H2HVL+H3H+RDCH
            DC    CLK2+$10000+PhRL+H1H+CEH+H2L+H2HVL+H3L+RDCH
SSKIP_END

; Assembler constant containing length of skip waveform table
SKIPLENGTH  EQU @CVI(SSKIP_END-SSKIP)
```

Care must be taken with the positioning of the waveform table in memory. In the above example, the table is 8 words long. The start address of the table must therefore begin at an address for which the lower 3 bits are all zero. In general the base address of the table must have the lower k bits =0 , where 2^k >= length of the table. If this condition is not met then the modulo address will be calculated incorrectly. This technique provides close to the maximum speed attainable with the clock board.

4.6. Image Header and Footer contents

The image header consists of the following 10 words:

L3 Status word
L3 gain setting (8 bits, 0=unity gain, $ff=max gain)
OP_MODE
OP_MODE
Frame count high
Frame count low
Integration time high
Integration time low
Number of rows per image
Number of columns per image

The image footer consists of the following 2 words:

VME	PCi
0	Frame Count high
0	Frame Count low

The OP_MODE, at start up, is as follows for the various applications so far implemented:

Mode 1 : $801
Mode 4 : $808
Mode 5 : $810
Mode 6 : $820

The L3 Status word at the start of the header contains the following information:

Bit 0 : Synthetic image mode set
Bit 1 : Pockels cell synch mode requested
Bit 2 : Pockels cell synch obtained
Bit 3 : The overexposure protection code has been triggered.
Bit 4 : We are less than 4 frames after a mode change or gain change

Bit 4 is used internally by the overexposure protection code in order to overcome bias transients that can cause false triggers. It has no use outside of the SDSU controller.

4.7. EEPROM burning.

The code for the L3 modes was all contained in an EEPROM on the SDSU timing card. Keeping the code in the EEPROM avoids any risk from spurious downloads that could destroy the CCD and makes mode changes very rapid. Burning this EEPROM is not straightforward. An S4 Dataman programmer connected to a PC serial port were used. The required EEPROM is the Atmel 28C256. The original NAOMI system used the 28HC256 for faster access but these are difficult to get and their extra speed is not really needed. Changing to the slower EEPROMS required some modification of the DSP initialization code that dealt with memory access wait states. This code is well commented and can be found in the boot source file timing_naomi_L3_initialisation.asm

The recipe is shown below, it must be followed to the letter:

```
1)    Program and burn all memory locations as FF
2)    Download bootcode  srecord file into Dataman as 'Motorola' format
3)    Move the bootcode in the Dataman memory :MOVE  8000-FFFF  -> 0000
4)    Burn bootcode from 0000-7FFF into EEPROM

5)    Now download application srecord files into the Dataman one at a time
      and burn code between the following locations only :

      App1  0-8FF
      App2  900-11ff
      App3  1200-1aff
      App4  1b00-23ff
      App5  2400-2cff
      App6  2d00-35ff
      App7  3600-3eff
```

The EEPROM has room for only 7 applications. Apps 4,5,6 are the Tip Tilt modes, App 1 is the initialization 'full frame CCD39' mode that MUST be present when used with the VME. The other application spaces are free for future upgrades, e.g. implementation of some of the more popular existing Shack Hartmann NGS modes.

4.8. Commands implemented in the WFS SDSU controller

The first group of commands are originals from the NAOMI system.

PON: Power on controller. Replies 'DON'. Also sets bit 12 of the X:<0 word (STATUS)

POF: Power off controller. Replies 'DON'. Also clears bit 12 of the X:<0 word (STATUS)

LDA <application number> : Load application from EEPROM into RAM. Application numbers 1-7. It needs to be followed by '**SYC 0 0'** to start the application running. Applications start with gain set to minimum i.e. unity L3 multiplication gain. This is done to protect against overexposure. Replies 'DON'.

SET <exposure time> : Set the exposure time, units are 25 microseconds. It needs to be followed by '**SYC 0 0'** to take effect. The CCD is a frame transfer device and so integrates the subsequent frame during readout of the current frame. The *actual* exposure time = frame readout time + requested exposure time. Replies 'DON'.

HIH: Set L3 gain to maximum. It needs to be followed by '**SYC 0 0'** to take effect. Setting gain hi in the first few frames after a mode change is not recommended as it could degrade the CCD, since it will be full of stray charge. Replies 'DON'.

SLW: Set L3 gain to minimum. It needs to be followed by '**SYC 0 0'** to take effect. Replies 'DON'.

ABT: Stop application running after current frame. Removes all clocks from CCD, clears synthetic image mode and any overexposure condition. Replies 'DON'.

TDL <param>: test data link when application is not running. Bounces back the original parameter. Useful to check the controller is attached.

RDM <address>: Read a memory location within the controller when application not running. Can be used to read an ID code in the EEPROM so that system knows which camera is attached. 'RDM $100006' should return the value $CCD60.

SYC <paramH paramL>: A hang-over from the original NAOMI code. It was used originally to synchronise the Master and Slave controllers. Some of the above commands need to be followed by SYC in order to function. The two parameters correspond to the high and low words of the frame count after which the previous command should take effect. For example 'SYC 0 50' would make the previous command take effect only after the frame count exceeded 50. The frame counter is set to 1 each time an application starts to run and is 28 bits long. Replies 'DON'.

The second group are all new commands that do ***not*** require '**SYC 0 0'** to follow. They are executed immediately after the current frame is read out and transmitted.

SSM: Set synch mode. The line readout becomes locked to the Pockels Cells . This will cause a slowing of readout since the controller waits at the start of each line for the Pockels cell to fire. In Mode1 the line readout time exceeds the Pockels cell interval (200us) so an artifact will probably be visible in a column towards one side of the image. The other Tip-Tilt modes have line times less than the Pockels cell intervals so should be clean of interference. Synch mode is cleared when the mode is changed. Pockels cell TTL trigger pulse needs to be connected to the BNC connector on the front of the SDSU controller. Replies 'DON'.

CSM: Clear Pockels synch mode. Replies 'DON'.

L3G <gain level> : set L3 gain level from 0-255, where 0 = slow speed gain (unity), 255 = Hi Speed gain. This command should not be used to set the gain high immediately after the CCD starts framing, since it will contain a lot of stray charge that will need a few frames to clear out. The exact values of the multiplication gain are defined in the EEPROM. Max gain is in the region of 200 although this value may experience drift with age/temperature. Replies 'DON'.

STI: set synthetic image mode. Images then contain a superimposed synthetic star image occupying the central 4 pixels of the image. The pixel values in the star image are 6000. Replies 'DON'.

CSI: Clear synthetic image mode. Replies 'DON'.

TTM <x-value> <y-value> : Send DAC values to the secondary Video board . This is included as an experiment to allow the SDSU controller to directly control the Tip-Tilt mirror. If a second Video board is included , configured with the links to be board #1 (board #0 is used for the CCD), then the x-value parameter will go to the DAC channel connected to 25 way D connector pin11, the y-value goes to pin12. These are 12 bit DACs with a range of +/-10V. Replies 'DON'. It was later possible to use a dedicated PCi DAC card to send signals to the TT mirror so this command was not really needed.

RLG : read back the L3 gain index. Initially used to cope with any server crashes that might make the TT PC forget the current gain setting. Gain setting later encoded in the header thus rendering this command obsolete.

RTM : read back the TT DAC values currently output by the SDSU secondary video board. Used to cope with any server crashes that might make the TT PC forget the current settings. Returns the X and Y DAC values.

5.0. SDSU Camera Controller

A standard SDSUII controller was used. An extra custom made multiplication clock board was fitted. Certain modifications were also required to the clock and video cards.

5.1. SDSU Clock Board

This needs clock shaping capacitors to be added as shown below. This gives the correct vertical clock slopes for reduction of Clock Induced Charge.

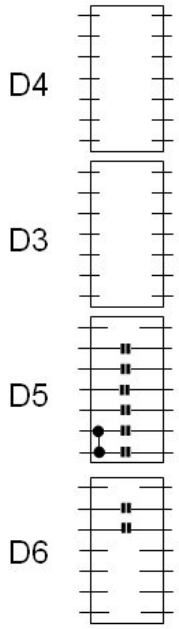

Positions of 2.2nF
Caps on L3 controller SDSU clock board.
Note link on D5 due to PCB track error

5.2. SDSU Video Board

This is modified as follows. CDS time constants are normally switchable between 1us and 5us with the default SDSU video card. To cope with the faster L3 readout the video card CDS capacitors are modified to give 500ns and 1us time constants. For this C3,C44,C74,C93 are all changed to 500pF. The L3 chips are operated with a non-zero substrate voltage provided by a DAC channel on the video board. DAC channels normally have a 1K series resistor. In the case of the substrate channel this load resistor (R143) must be replaced by a zero ohm link.

5.3 SDSU Timing Board

This requires the L3 WFS EEPROM to be fitted. The hardware remains unmodified, however, for the Pockels synchronization to work, a link (JP19) must be fitted to the board. Additionally, JP18 must be absent. Their positions are shown in the picture below.

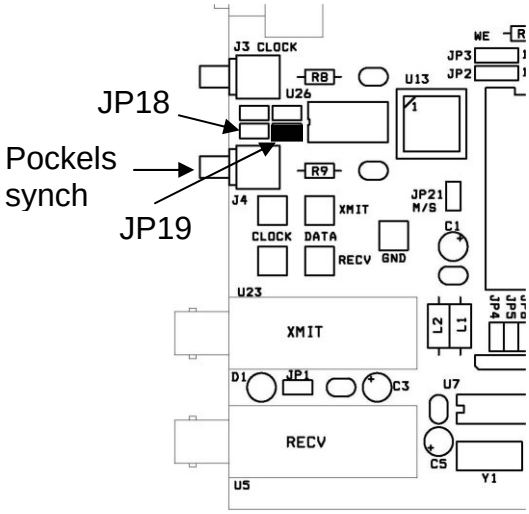

5.4 SDSU Utility Board

This is not required by the WFS, however, it is included in the controller so retain spares compatibility with other L3 science cameras.

5.5. ING L3 Multiplication Clock Board Design

The L3 CCDs require a high voltage clock to be applied to the H2HV pin to induce avalanche multiplication. Avalanche gain commences when the amplitude of this clock reaches approx 35V. The gain is then a steeply rising function of clock amplitude. 50V is the absolute maximum amplitude, beyond which damage can occur to the CCD. The standard SDSU controller can only supply a 10V clock so a new clock driver was built in house and fitted into a spare SDSU slot. A circuit diagram can be found in the appendices. Here the circuit is described in more detail.

5.5.1. Design Description

The card had the same dimensions as a normal SDSU card. It attached to the backplane in the normal way. No signals are obtained from the backplane, only 5V and +/-16V power as well as ground. Signals are fed to and from the card via a 9 way D connector on the front of the card. There is also the option to make the ground connection through this connector instead of the backplane by disconnecting link J10 and making link J7. This was included in the design in case any earthing problems/ground loops were encountered.

5.5.2. 60V Supply

Two +/-15V DC to DC converters connected in series are used to generate a 60V high voltage bus. No single 60V unit was available.

5.5.3. Programmable Voltage Regulator

The high voltage bus feeds into a regulator comprising U1,Q1 and Q2. A rising output on op-amp U1 robs base current from the power transistor Q1 and reduces the regulator output voltage. Zener diodes D3 and D4 together dictate the maximum output voltage from this regulator i.e. when the output of U1 is zero.

Two further op-amps (U2 and U3) as well as a voltage reference (U4) are included in the circuit to allow the regulator output to be set by one of the DACs on the SDSU video board. The resistors associated with these op-amps are all low temp co-efficient varieties.
The video board DAC range is +/-5V. When set to zero volts (or when disconnected) the voltage regulator drops to approximately 20V output. When the DAC is set to its maximum positive value the regulator reaches its maximum voltage. The negative part of the DAC range is not utilised and diode D1 ensures that –ve DAC voltages are in fact clipped to zero.

5.5.4. Output switching circuit

The voltage regulator output is connected to the power rail of the output switching stage via a passive filter consisting of an inductor and an RC smoothing network (L2,R7,C9,C20). The values of R7 and C9 are critical to avoid very unpleasant power on/off transients at the clock output. If C9 is made too large then the inrush current at power on is sufficient to trip the power control board in the SDSU. Inrush current can still be a problem if the circuit is powered on with the DAC voltage set to maximum. It is recommended that the SDSU code sets the DAC voltage to zero during the power on sequence, increasing to the higher L3 gain levels a few hundred ms after the controller is fully powered on. The WFS code does indeed do this, and starts up in low gain mode requiring the user to later switch the gain to an operational level.

The output drive transistors are biased using Schottky diodes D5,6,7,8. These transistors are driven through RC shaping networks comprising C12,13,14,15 and R16,17. These values can be trimmed to vary the shape of the edges of the output waveform. The input to this switching circuit is obtained from a spare channel on the standard SDSU clock card and is buffered on board using two 74HCT04 inverters connected in parallel. The push pull output passes through a 100 Ohm resistor to reduce ringing. An optional pulse shaping capacitor can be connected to J6 to smooth the edges of the clock waveform although in practice this was not utilised as the clock edges could be controlled sufficiently using the shaping networks attached to the transistor bases.

ING HV board mounted in an SDSU II controller

5.5.5. Tuning the performance

Resistor R14 was chosen with care. It dictates the control range of the DAC i.e. what the output voltage of the clock generator will be when the DAC has its maximum output. It is very important that R14 is not too small or else the transistor Q2 will become saturated when the DAC is below its maximum voltage. This means that when high voltages are demanded the top part of the range will not be fully regulated. R14 was chosen by first setting the DAC to its maximum and then reducing R14 steadily until the Vcb of Q2 reached approx 1V.

The gain of an L3 chip is critically dependant on the phase of the H2HV clock. In particular, the H2HV clock must be at its maximum value at the moment that clock H1 goes low on the CCD. Small phasing errors between these two clocks can give huge changes in gain. For this reason three extra test points TP9, TP10 and TP11 are included at the edge of the board. These three are undedicated, and simply connect directly to pins on the 9 way edge connector. A further test point TP7 at the board edge is connected to H2HV. These undedicated test points can be used to scope the other clock phases of the CCD i.e. those generated normally by the SDSU clock card. Having all the test points close together like this is much more convenient.

6. Special features of Camera

Some extra software features were added to the camera to improve its reliability.

6.1. Overexposure Protection

Saturating levels of exposure (i.e. full well) combined with high L3 gain can damage and even destroy the CCD60 if sustained for long periods. The exact relationships between onset of damage and level of overexposure are not fully known. The SDSU code therefore contains a highly conservative ‘gain killing’ subroutine that looks for overexposure in the pixel data and drops the L3 gain to 1 if detected. Pixels in the penultimate row of the image window are analysed and if the average value is >3000 ADU above bias, the gain is dropped back to 1 within about 50ms. Subsequent images as displayed on the TopGui show a small black rectangle in the bottom row to indicate what has happened. When using the tip-tilt GUI, the overexposure condition is flagged by a red-text alert at the top of the window. A bit is also set in the OP_MODE word of the image header. If the gain index is already set to 0, the overexposure protection is not needed and is disabled. Once overexposure has occurred the state can only be cleared by a mode change (or reload of the current mode) or a manual change of the gain level or readout speed. This is a small inconvenience to the user but a necessary one for the security of the detector.

Although a stellar image exposes linearly up to about 20KADU on the WFS, a flat field over-illumination such as that may be obtained from flooding with room lights gives a much lower level of saturation : around 5K ADU.

6.2. Pockels cell Synchronisation

The WFS SDSU code provides the option of synchronizing the line readout of the CCD to the Pockels cell opening. This cell opens for approx 2μs every 200μs. It uses a high voltage pulse to

open and it seems likely that some form of interference will result. The line readout for the 8x8, 16x16 and 32x32 pixel Tip Tilt modes is less than 200μs so adding a small delay routine at the start of each line that waits for the Pockels cell to fire can ensure that the EMI does not appear in mid-line and cause an image artifact. The Pockels cell is triggered by a TTL signal from a timing rack. This signal must be tapped and connected to the BNC connector on the front of the WFS controller. A negative edge on this line triggers an interrupt in the SDSU DSP. This interrupt sets a status bit that is continually read by the above mentioned delay routine. If no synch pulse is detected within a fixed time interval, the loop exits and the next line is read anyway.

The 88x80 pixel Mode 1 image format takes approximately 300us per line-read. When synched to the Pockels cell, the Pockels pulse should therefore still be visible about 2/3rds of the way across the image.

Synched mode will slow the readout time of the WFS by between 1.1ms and 1.7ms, depending on the mode. The synch pulse from the front panel connector goes to J4 at the edge of the timing board. Even slower readout will result if Pockels synch mode is set but synchronization not achieved, for example if the synch cable is not fitted.

7. Shortcomings and future improvements

The readout speed is limited by the time it takes to skip over unwanted pixels outside of the image window. This is done as fast as the clock driver board will allow i.e. with 40ns overlaps between the serial clocks. This is far short of what the CCD is capable of. A medium sized project for the future would be to extend the functionality of the HV clock board so that it generates all the serial clocks. It could contain a small CPLD based sequencer dedicated to the CCD60 that functions at a higher speed than the standard ARC clock board.

The HV clock board had a poorly regulated HV bus voltage that dipped slightly when the clocks were active. Due to the high sensitivity of the L3 gain process on multiplication voltage, this had the tendency to increase gain massively at the start of readout, resulting in image artifacts. The solution was to fully exercise the clocks even during frame transfer and integration phases. The bus instability is in fact due to the RLC network on the clock card used to prevent inrush currents at start up.

Further experimentation should be done with the parallel clock drivers to try and reduce the level of Clock Induced Charge.

Further cooling could be produced by using larger gauge cables to carry the Peltier current and/or boosting slightly the Peltier supply. The dark current is, however, insignificant even with the existing system. The main motivation for further cooling is to boost the multiplication gain yet further either if the camera suffers excessively from pickup noise in the GRACE environment or if the gain mechanism undergoes ageing.

The intrinsic read noise at 54e is very high although its contribution is all but eliminated by the internal multiplication gain. It should be in the region of 15e. This poor performance is probably

due to the fact that the camera head is electrically shorted to the rest of the WFS structure and it would have been better to use an electrically isolating mount. Some pickup is also seen from the vertical clocks of the Master WFS, which is close by. With an L3 chip these effects can be further removed by increasing the multiplication gain.

A major part of the readout cycle is simply skipping unwanted pixels. It may be possible to clear pixels to the right of the read out window using the on-chip Dump Gate. This would take just a few microseconds for the whole row rather than about 0.6μs for each pixel skipped. This was briefly experimented with but its effects on the pipeline were very complex. This should be revisited since it could make a big difference to read-out speed. If it could be made to work, further improvements would then come from placing the guide star window at low CCD column numbers.

Appendices

Wiring of SDSU controller, WFS head and interconnecting cable

Additional Wiring added to L3 WFS :

Synchronisation connector
Coaxial cable taken from J4 connector at edge of timing board to a coaxial connector on the front panel of the controller. This allows synchronization with the Pockels cell.

Tip Tilt Mirror connections
Experimental connectors to allow the SDSU to directly drive the TT mirror. Two coax connectors added to the front panel, going to pins 11 and 12 of the secondary video board. Not used in final system and the valuable video board was removed for use elsewhere

L3 WFS Wiring

Function	CCD60 device pin	CCD PCB 21 Way uD	WFS Head 26 way	Front Panel 18-32	ING HV Board (9wayD)	SDSU Clock Board (37wayD)	SDSU Video Board (25WayD)	
IPhi1	6	purple upper	A	A		13		
IPhi2	29	grey lower	B	B		14		
				C		15		
SPhi1	3,30,28	yellow upper	C	D		16		
SPhi2	2,27	orange upper	D	E		17		
				F		18		
DG	26	red upper	E	G		19		
PhiR	10	green lower	F	H		1		
				J		2		
RPhi1	22,24	red lower	Z	K		3		
RPhi2	21,25	orange lower	a	L		4		
RPhi3	23	brown lower	b	M		5		
RPhiDC	14	brown upper	c	N		6		
RPhiHV	15	black upper right	L	P	6			coax core
Sheath-RPhiHV				d	7			coax sheath
spare clock				R		7		
spare bias				S			7	
VSS	7,11	whiteL,blackL,whiteU,greyU	M	T			12	
OS	12	yellow lower	N	U			SMA1 core	coax core
Sheath-OS(L3)				f			SMA1 sheath	coax sheath
				V			SMA2 core	coax core
				W			SMA2 sheath	coax sheath
ABD	4	green upper	P	X			5	
				Y			6	
IG	5	blue upper	R	Z			11	
RD	19	purple lower	S	a			3	
				b			4	
OD	13	blue lower	T	c			1	
				e			2	
OG	20	black upper left	U	g			9	
Camera Chassis			V	h		26		
L3 clock enable					2	8		
L3 Clock input					3	9		
L3 GAIN					4		10	
Diagnostic clock				j		20		
					7	24	13	>>to Star Point

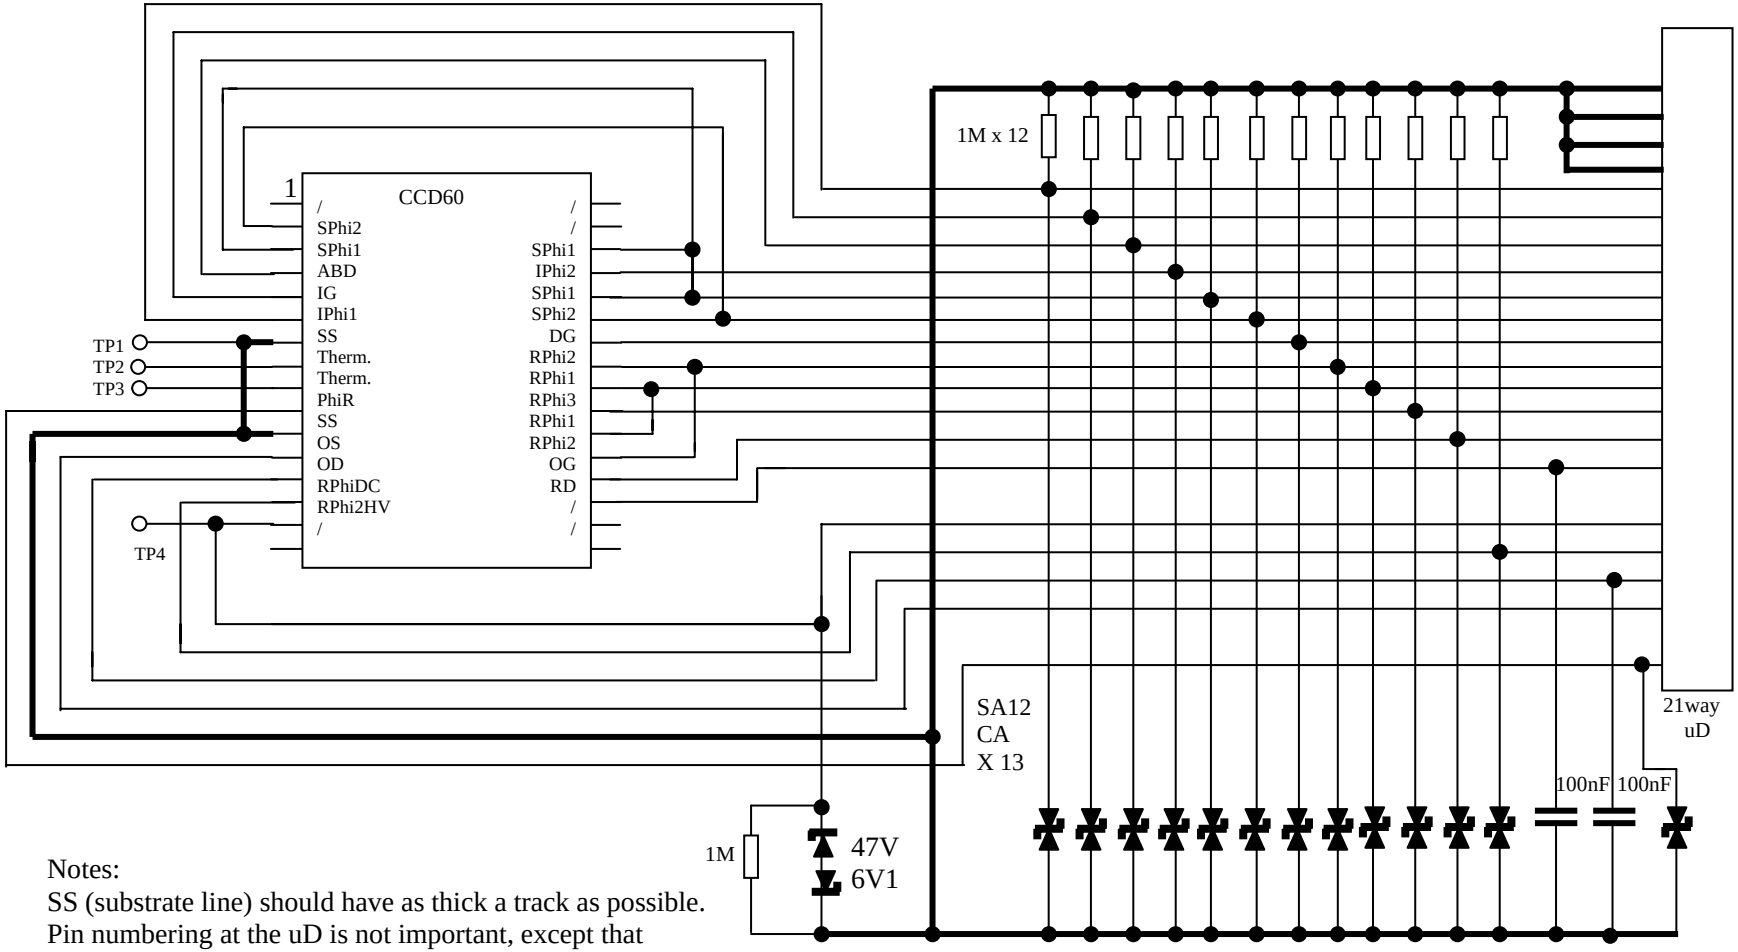

Notes:
SS (substrate line) should have as thick a track as possible.
Pin numbering at the uD is not important, except that
RPhi2HV should be placed at one corner isolated from
the other connections by nearby SS connections to minimise pickup
On other signal lines.
Static protection diodes same as those used in AS box. The anti-static diode on PhiR was unfortunately absent from the PCB design and was added manually later.

CCD60 WFS Head Board

SMT July 04

SDSU ->CCD60 WFS Head Cabling

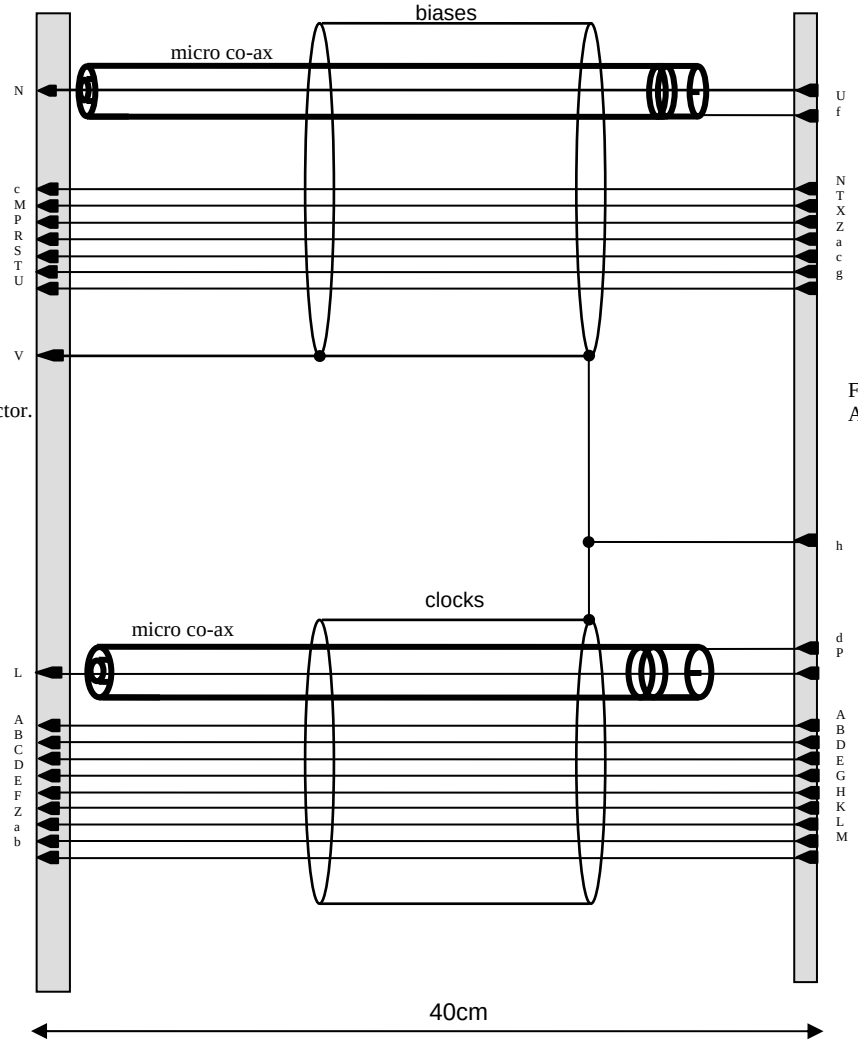

CCD60 WFS Head Connector.
Amphenol 16-26 Hembra

Front panel connector.
Amphenol 18-32 Macho

Simon Tulloch Oct 14 2004

L3 WFS SDSU Controller Internal Cabling

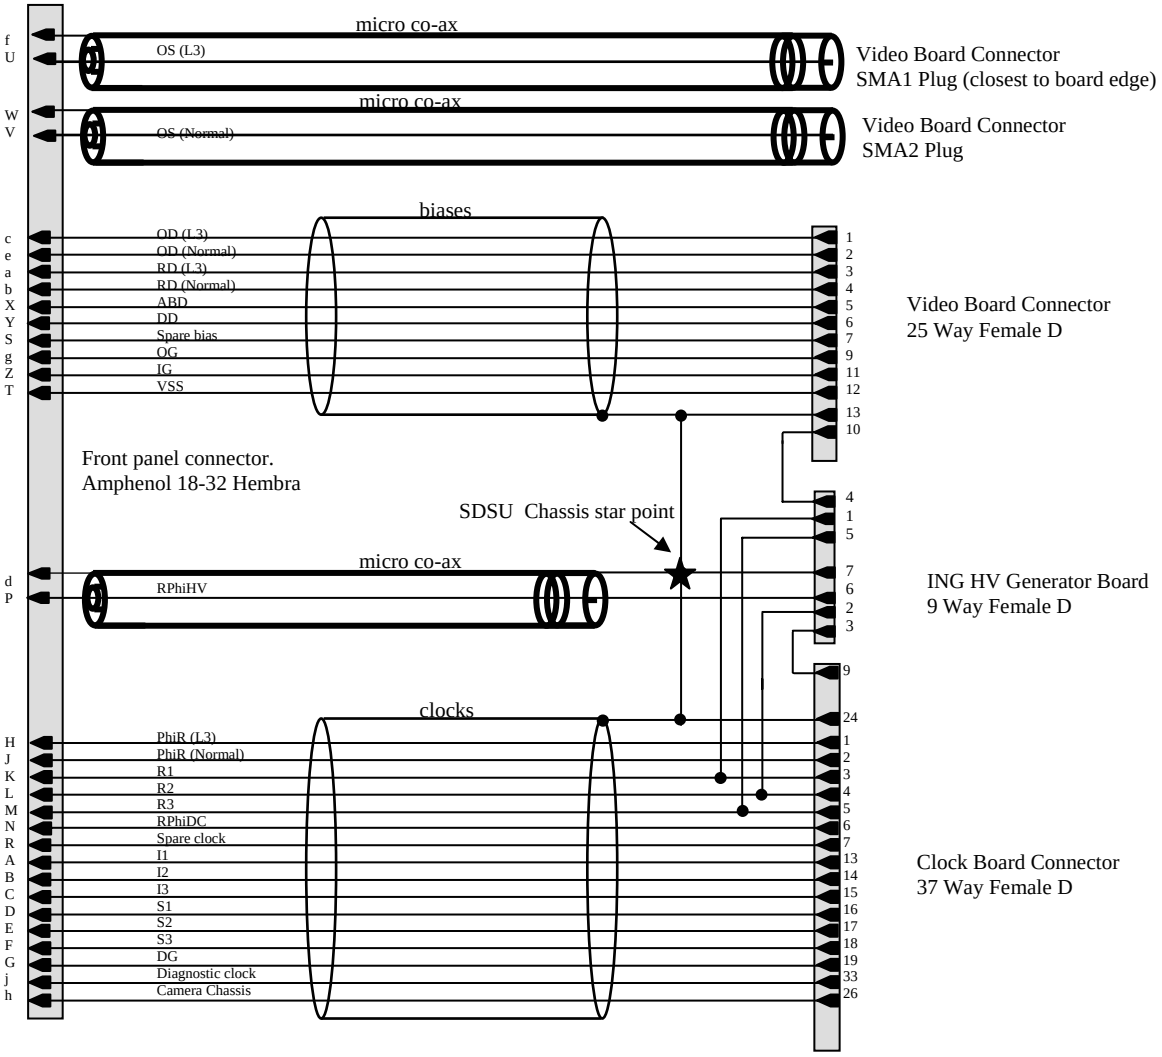

Simon Tulloch Oct 14 2004

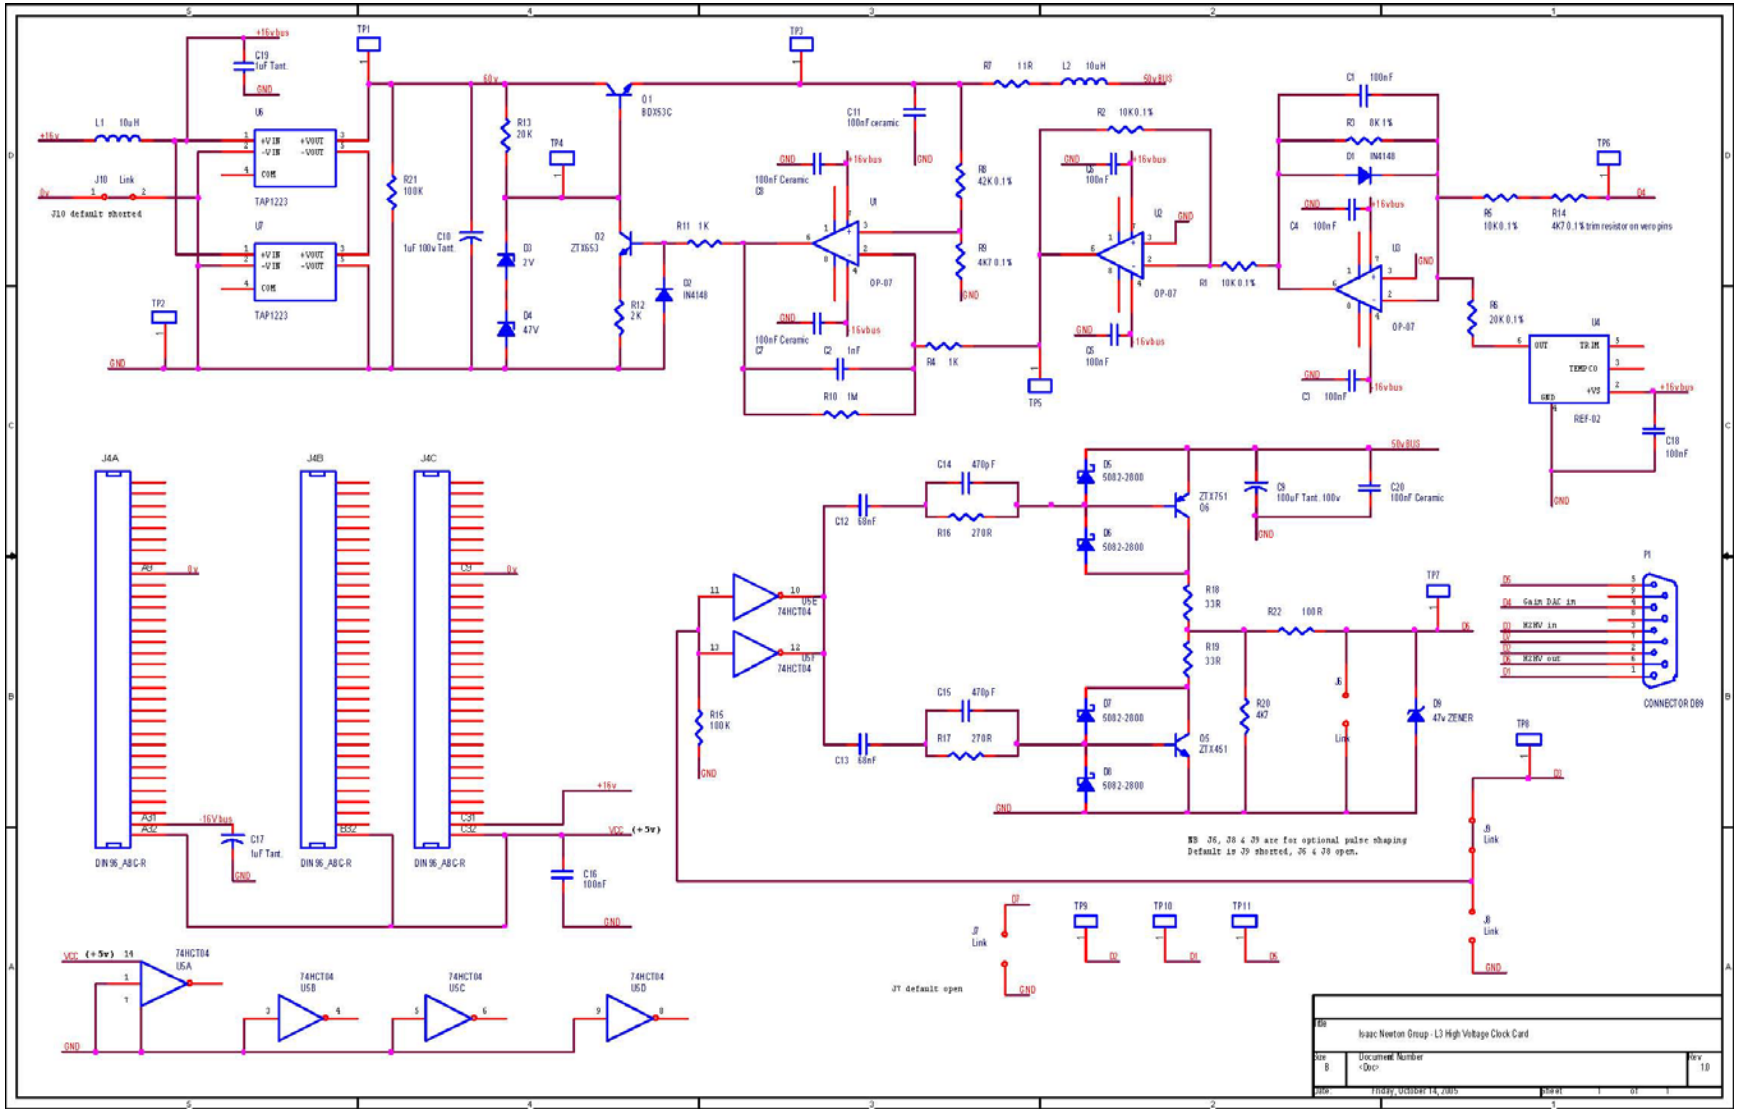

ING L3 HV Clock Generator Board