

Hawkeye Ocean Color Instrument – Performance Summary

Alan Holmes, John M Morrison²,

Gene Feldman³, Fred Patt⁴, Shihyan Lee⁴

¹ Cloudland Instruments, Inc, ² University of North Carolina Wilmington (UNCW)

³ NASA Goddard Space Flight Center (NASA),

⁴ Science Applications International

Cloudland Instruments, 1124 N Fairview Ave, Goleta, Ca 93117

+1 (805) 403-6620; alan@cloudlandinstruments.com

ABSTRACT

Hawkeye is an ocean color instrument that is part of the SeaHawk satellite developed for SOCON, the **Sustained Ocean Color Observations using Nanosatellites** program funded by the Gordon and Betty Moore Foundation and managed by the University of North Carolina – Wilmington (UNC-W). HawkEye has spectral characteristics similar to SeaWiFS, but with 8 times finer resolution and a smaller field of view more appropriate for lakes, rivers, and near-shore terrestrial environments. With a volume of only 10 X 10 X 10 cm (a CubeSat 1U), it can produce 8 bands of image data in a single pass, each with 1800 x 6000 pixels, with a resolution of 120 meters per pixel. This paper will present a short summary of instrument design, the spacecraft interface, and “lessons learned” during this effort. Scientists considering using linear arrays in a pushbroom mode for remote sensing will find this useful. Much of the discussion will center on optical performance, such as flat field calibration, polarization effects, stray light, out-of-band response, and exposure linearity. Images from field tests will be shown.

Keywords: SEAWiFS, ocean color, Hawkeye, Seahawk, Gordon and Betty Moore Foundation, UNC-W, Clydespace


1) MISSION SUMMARY

The Hawkeye instrument is an ocean color measuring instrument designed to fly on the SeaHawk satellite developed for SOCON, the **Sustained Ocean Color Observations using Nanosatellites** program funded by the Gordon and Betty Moore Foundation and managed by the University of North Carolina – Wilmington (UNC-W). The Hawkeye instrument measures ocean color in 8 spectral bands, similar to SeaWiFS, except Band 7, which is shifted to a slightly lower wavelength to avoid the oxygen absorption feature that a wider band overlapped on SeaWiFS. The instrument is approximately 1/3rd the volume of the entire satellite, which is a 3U Cubesat manufactured by Clydespace in Glasgow, Scotland. The purpose of this instrument is to ascertain the quality of ocean color data possible with such a small, inexpensive instrument and bus. The nominal orbit is 540 km, and the nominal pixel geometric instantaneous field of view (GIFOV) 120 meters on a side. Each band will produce an image 1800 x 6000 pixels in size, for a total field of view of 216 X 720 km.

2) DESIGN CONCEPT


The Hawkeye instrument uses linear arrays in pushbroom mode to collect data over a two dimensional area. The instrument has 4 linear CCD arrays, the Onsemi KLI-4104, to collect the 8 bands of data. Figure 1 illustrates the optical design for two bands, sharing a single array.

Figure 1: Two Bands share a Single Linear Array




Each band has its own optical system. Figure 1 shows 2 of the 8 triplet lenses used in this instrument. Each triplet lens is preceded by a polarization scrambler to minimize polarization variation of response, and passband filter, and followed by multiple baffles to isolate the bands. The instrument contains 8 bands, arranged as shown in Figure 2.

Figure 2: The Instrument contains 8 Bands



Four KLI-4104 linear CCDs are used. Each one has three Chroma channels, which in a document scanning application would be covered by a Red, Green, and Blue filter matrix. Figure 3 illustrates the CCD pixel geometry.


Figure 3: CCD Geometry



In our implementation the Chroma channels are monochrome – the filter masks are omitted. We do not use the Luma channel incorporated in this CCD. Our triplet lens is F/5, which provides enough light to the detector that we can readout the CCD 4 times per pixel dwell time. So, for each pixel on the ground we collect 12 pixels worth of data – 4 looks times 3 Chroma channels, substantially lowering the random noise in the final image, and smoothing dust or pixel level nonuniformities. Since the three chroma channels are physically offset from one another we have to shift the data from them appropriately before summing.

The four CCDs are contained on one CCD board. The lenses are translated to register and focus each band to the others. Underneath the CCD board we have an Analog board, which contains all the high current clock driver and level shifter electronics, 12 low noise preamps, and four 16 bit multichannel A/D converters (AD9826) implementing a correlated double sampling readout. Current, voltage and temperature telemetry are collected using an AD7490 16 channel SPI-based A/D converter. A mother board assembly is used to pass these signals on to a PicoZED board used as the instrument computer. Figure 4 shows the arrangement of the circuit boards in the instrument.

Figure 4: 6 Circuit Boards are used in Hawkeye



The PicoZED board contains a ZYNQ 7Z015 ARM dual core processor with 1 Gigabyte of DRAM, 148 pins of user I/O (which we needed), and 128 Mbytes of QSPI flash for booting the FPGA. It basically integrates an FPGA with a LINUX computer. This processor was chosen for its speed, since our design has a great deal of data to process on the fly. The FPGA implements a readout engine for the finderscope CMOS array, and for the 12 channels of linear CCD data, with support for 1:1, 2:1, and 4:1 oversampling. The data collected is transferred to the Linux side of the processor using DMA. The Linux side uses Digilent Linux version 3.3. The Linux program implements spacecraft command

and control over an RS-422 link to the spacecraft, receives the array data from the FPGA, post processes that image data to shift and aggregate the multiple channels of data for each band, packs it for transmission, and then sends it back to the FPGA for transmission to the spacecraft computer over a Quad SPI download port at 5 MegaBytes per second. Ethernet support is included for debug and development, but not brought out to the spacecraft.

Considerable effort went into designing the internal instrument power supplies, which were required to produce multiple voltages for the processor and CCDs. The instrument overall draws 1.45 amps of current at 7.2 volts during nominal operation. However, during activation, it was necessary to limit the inrush current pulled by the instrument to be no greater than 3 amps. Normally one assumes with battery operation that momentary high currents are not a problem, but the Cubesat batteries have a high enough worst case internal resistance that a high current can draw the bus voltage down to where the spacecraft computer resets. This was a non-trivial problem, solved by slowing down the internal supply transients during initiation.

Figure 4 also shows two solenoids which are used to drive a thin mechanical shutter back and forth in front of the linear CCD arrays, a capability which allows dark data to be collected at the end of each frame of pushbroom-scanned data. This data is used to remove the thermal noise due to hot pixels from the image. The hot pixels are not a large effect here for the short exposure times used but would add to the noise level if not corrected. Up to 50 pixels can be averaged for each column to calculate the dark offset, so the dark data does not contribute to the readout noise. A “finderscope” assembly is also shown. The finderscope camera captures wide field images sequentially during an imaging pass, which are saved and can be used to determine any pitch, yaw or roll that occurs during imaging to higher accuracy than the spacecraft attitude control data allows. The finderscope is filtered to have near infrared response around 850 nm to show clouds and land features with maximum contrast against the dark ocean. Each finderscope image is 752x480 pixels in size, and up to 25 images can be captured during a pass. The instantaneous field of view of the finderscope is comparable in width to that of the linear CCDs.

The final unit is covered by an aluminum skin to achieve a light-tight enclosure in orbit. It is easy to forget that these instruments are designed to be used outdoors in full sun, staring at a bright cloud covered earth, when one is working in a darkened laboratory. A completed instrument is shown in Figure 5.

Figure 5: Hawkeye Unit One



Challenges of Ocean Color Imaging and Cubesat Operation: a Cubesat is a very small platform, with a final weight in the range of 6 kilograms. Attitude control, accomplished through sun sensors and reaction wheels, is a challenge, since the Hawkeye resolution is on the order of $1/80^{\text{th}}$ of a degree. Figure 6 shows an instrument model (gold-colored) in the Cubesat frame, to give an idea of the size of the complete satellite.

Figure 6: Instrument in Cubesat Frame




Figure 7 shows a picture of the full moon captured with the instrument for illustration of the resolution.

Figure 7: Full Moon imaged by Engineering Model Hawkeye



The resolution is comparable to the human eye. Clydespace has determined that the drift in attitude will be about 1 degree rms in all axes, but with a long time constant, so it is hoped that this smooth behavior, augmented by our finderscope camera, will allow full resolution to be achieved. Another issue is that a free-floating Cubesat is a compass needle, and activation of our solenoids will produce a significant perturbing magnetic torque. That is why the shutter is normally open, and the solenoids are only activated at the end of an exposure. As a quantitative example, if we left a solenoid on for the duration of the exposure, 100 seconds, it would rotate the Cubesat about 5 degrees out of position. We use two solenoids, wired in opposition, so their residual fields mostly cancel out. The spacecraft torque rods are powerful enough to deal with the residual torque. These effects are hard to calculate, but should be measured early on in any program by hanging the instrument from a string and using torsional pendulum calculations to determine the magnitude of the magnetic dipole interaction with the earth's field.

3) CALIBRATION AND ARTIFACTS

3.1 Achieving good signal to noise ratios (SNRs) in the radiometric data:

Ocean color data collected from orbit is corrupted by the reflection of light from the atmosphere and the reflection of that light off the ocean's surface – the ocean itself is quite dark. 90% of the light sensed is not from the ocean in some bands. Therefore, an instrument has to have good resolution of the light levels so that useful data is available after removal of the atmospheric component. This is why we employ the 4X oversampling coupled with the aggregation of

3 looks (the Chroma channels) to improve the SNR by the square root of 12. Table 1 shows the expected SNR expected for typical ocean radiance levels. This actual performance matches this theoretical performance closely. We achieve a readout noise of 25 electrons rms with our design, a level that is much less than the photon noise inherent in all bands.

Table 1: Expected SNR Performance for Typical Ocean Radiances

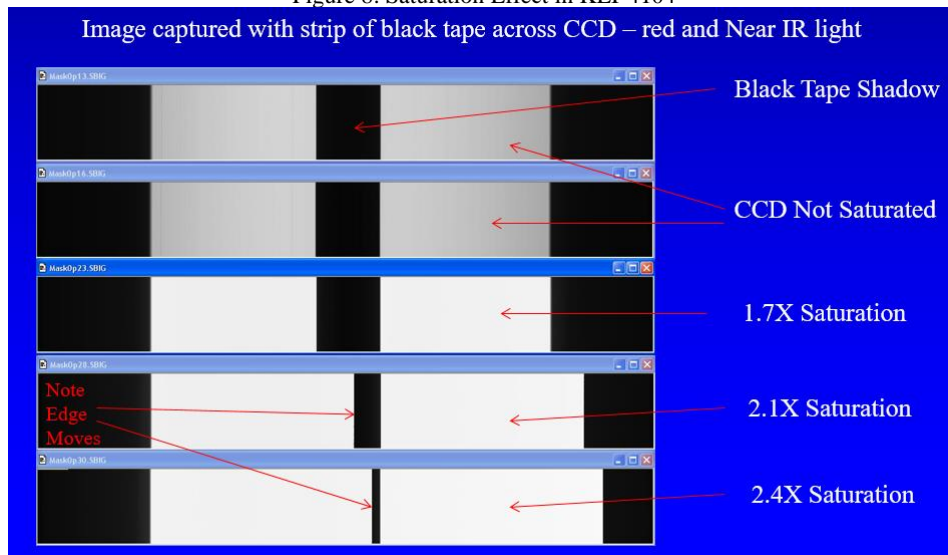
Standard (Ltypical) Radiance			Radiance					CubeSat
Band	Wavelength	Bandwidth	Watts/m ² /ster/micron	Optical	QE	Photons	Signal in	SNR per
	in nm	in nm		Trans		per watt	Electrons	pixel
SeaWiFS 1	412	19.31	78.6	0.663	0.59	2.07E+18	11523	362
SeaWiFS 2	443	20.45	70.2	0.663	0.64	2.23E+18	12823	383
SeaWiFS 3	490	19.84	53.1	0.774	0.72	2.47E+18	13726	397
SeaWiFS 4	510	19.15	45.8	0.760	0.75	2.57E+18	12109	372
SeaWiFS 5	555	20.73	33.9	0.682	0.74	2.79E+18	9359	324
SeaWiFS 6	670	20.58	16	0.653	0.68	3.37E+18	4632	221
New 7	750.9	15.61	9.3	0.818	0.49	3.78E+18	2072	138
SeaWiFS 8	865	39.06	4.5	0.843	0.28	4.35E+18	1696	122

The SNR can be further improved by binning the final image. With 2 x 2 binning the image is still 900x 3000 pixels in size, but with SNR performance 2X better, comparable to SeaWiFS for most bands, and with 4 times the resolution.

3.2 Controlling CCD Blooming Effects:

CCDs are very good detectors but have a few artifacts that must be considered. First of all, many CCDs “bloom” when overexposed. This can be an issue in ocean color since the temptation is to design for the light levels typical for the ocean, and ignore the saturation of bright clouds. However, tropical regions tend to have puffy white clouds, and also coastlines can have bright regions of rock or sand near the coast, or jungle foliage which is quite bright in the near infrared. If these regions saturate the detector they can bloom into the clear areas quite badly. The KLI-4104 CCD used here is pretty good in this regard. Figure 8 illustrates the CCD response to saturating light levels. For this figure a strip of black electrical tape was placed across the CCD to mask it. Note that as the light level increases the saturated areas expand to the right.

Figure 8: Saturation Effect in KLI-4104



We have determined that this blooming is insignificant at exposure levels up to 1.7X saturation, at which point it rapidly expands, so we have designed the nominal exposure and optical throughput so that high bright clouds are no more than 1.7X saturation in the image. This was the design exposure value used in Table 1.

In any new design that might use Onsemi area arrays based on older Kodak designs, a designer should be aware that full frame CCDs tend to bleed badly, but interline parts are quite good, in general.

3.3 Peculiar CCD behavior with Short Exposure Times:


The instrument samples a line of data at a spacing set by the interval time. This time is set by the orbital altitude and ground speed. Within each interval time 4 exposures are captured with a duration set by the exposure time. The exposure time can be less than $1/4^{\text{th}}$ of the interval time to reduce the amount of over-saturation from clouds, for each band. The KLI-4104 CCD has a line one can use to holdoff the beginning of the exposure in a measurement cycle, which provides exposure adjustment while keeping the row rate constant. During testing we found an anomalous behavior with short exposure times for the KLI-4104 CCD. The 12 inch traveling NASA sphere we used for calibration is quite bright in the near infrared. We found that, at red and near infrared wavelengths charge is leaking past the holdoff switch, so the net exposure is a combination of the leakage over the interval time and the intended exposure. The effect is very deterministic but needs to be understood for accurate knowledge of the “exposure time.” Table 2 below gives the relative magnitude of the leakage to the intended exposure. One can use this information to calculate the actual exposure from detailed knowledge of the CCD exposure timing. Since the row rate may be much longer than the exposure duration this effect can be significant, more than the 8% shown for Band 8.

Table 2: Leakage Path Fractional Magnitude

Summary of Both Units	Leakage	Leakage
	Fraction,	Fraction,
	Unit One	Unit Two
Wavelength		
412	0.0023	0.0039
443	0.0045	0.0056
490	0.0096	0.0099
510	0.0132	0.0137
555	0.0144	0.0156
670	0.0415	0.0433
751	0.0630	0.0611
865	0.0863	0.0808

This effect has huge implications for flat fields. We use the same sphere for collecting flat field data (where the optical system is flood illuminated) to calibrate out the pixel-to-pixel variations caused by CCD nonuniformity and dust. Unfortunately, the leakage path has an entirely different flat field pattern from the direct path. If flat fields are captured with different combinations of exposure and interval time the components can be mathematically separated. Figure 9 shows the two flat field components on the same scale. Note that the leakage path has greater non-uniformity than the direct path, so the fact that it is only 8% of the final total does not allow one to ignore this contribution.


Figure 9: Flat Field Differences between Leakage and Direct Path



3.4 CCD Temperature Effects:

The KLI-4104 CCD has a significant variation in response with temperature. This was measured during Thermal Vacuum calibration testing and is shown below in Figure 10. The anomalous point at minus 10 degrees C for band 1 should be discounted since it was not seen in Unit One data and is probably due to some slight movement of the chamber relative to the table with the sphere over the hours of measurement. Also, the temperature effect is much greater in the near infrared for both units.


Figure 10: Temperature Effect for Unit Two



3.5 Suppressing Out-Of-Band Response:

Out-of-Band leakage refers to a common problem with any design in that filters may allow radiation far outside the passband to leak through at the 1% or even 0.1 percent level. While this doesn't sound too bad, the problem with a narrow band filter instrument is that a much wider passband may be contributing to the leak, and the detector may have better response to the leaked radiation than the desired signal. A good example is a system that uses a silicon photodiode to measure blue light. The silicon photodiode's excellent response to near infrared light will enhance any leakage. The filters in Hawkeye were procured from Omega Optical in Brattleboro, Vermont and are quite good in this regard. Figure 11 shows the out of band response for every band.

Figure 11: Hawkeye Filter Transmissions



One must multiply this transmission by the source spectral distribution and the detector quantum efficiency to determine the percent of out-of-band leakage in each band. We have done this and the result is shown in Table 3. The out-of-band percentage is tabulated for a solar spectral distribution, as well as an ocean (Ltypical) spectrum, and the traveling 12 inch NASA sphere we used for calibration. Integrating spheres tend to be very red, so large calibration errors can result from this effect if one has a red leak.


Table 3: Out-of-Band Prediction based on Actual Filter Data

Band	1	2	3	4	5	6	7	8	
Center Wavelength in nm	411.3	446.2	487.0	509.6	554.4	669.4	749.8	863.9	
Peak Transmission	69.8	69.8	90.3	88.6	91.1	87.2	86.1	88.7	(Percentage)
Bandpass in nm	18.8	22.0	19.8	20.2	20.3	20.4	16.2	39.6	
Out-Of-Band - Solar Spec	0.6	0.8	0.9	1.0	1.2	1.4	1.4	2.4	(Percentage)
Out-Of-Band - Ocean Spec	0.5	0.9	1.1	1.2	1.8	2.3	1.6	4.0	(Percentage)
Out-Of-Band - Sphere Spec	3.0	1.2	0.8	1.0	0.9	1.2	1.3	1.6	(Percentage)

3.6 Minimizing Polarization Dependence of Response:

Light scattered from the earth's atmosphere can exhibit strong polarization. It is important that an ocean color instrument not be sensitive to polarization to produce an accurate measurement. We designed the instrument to have minimal polarization sensitivity. However, early measurements showed that the CCD and triplet lens combination, with no filter or scrambler, had about a 3% degree of polarization on-axis. For this reason we added a single wedge polarization scrambler into the optical path for each lens with the crystal axis of the scrambler at 45 degrees to the linear CCD. During acceptance testing of the full system we found that the scrambler did indeed reduce on-axis polarization sensitivity to zero, but our passband-defining filters had introduced significant off-axis sensitivity. Figure 12 shows the response across the Field of Regard as a polarizer is rotated in 15 degree increments between the finished instrument and an integrating sphere. Each 15 degree measurement is separated by a black line.

Figure 12: Unit Two Band 6 Response with Polarization



As you can see, a two-cycle variation in response over a 360 degree rotation is occurring, indicative of polarization sensitivity. The degree of polarization varies with passband, as shown in Table 4.

Table 4: Hawkeye Instrument degree of Polarization

	Average	Average
Band	Unit 1	Unit 2
1	1.22	1.69
2	3.26	3.78
3	3.48	3.07
4	3.27	3.12
5	1.66	0.99
6	4.11	3.78
7	4.71	3.88
8	1.16	1.36

What the measurements shows is that we have maximum polarization sensitivity at a 45 degree polarizer angle where the scrambler has no effectiveness. Component testing indicates this polarization sensitivity is caused by off-axis light passing though the filter. The filter is not used at more than 12 degrees off-axis, so this was surprising. Had we not used a scrambler at all this variation would have been several times worse! In retrospect we should have used a double-wedged scrambler to further reduce this sensitivity. Beware that scramblers introduce some blur since the two polarization states are refracted by different amounts by the scrambler's component prisms. Here we chose a wedge that resulted in one pixel of blur, barely noticeable.

3.7 Eliminating Back Reflections to the CCD:


In an optical system design the natural tendency of the designer is to make all components square to the optical axis. This can lead to very bad stray light problems if the lens system has a filter in front of it, as Hawkeye does. The KLI-4104 CCD silicon substrate visually looks like a very shiny 1 mm wide mirror with a few barely visible black lines

corresponding to the pixel active areas. Light from the target scene that strikes this substrate is specularly reflected right back out of the system. The light strikes the filter, and is partially reflected back into the optics where it travels to the detector again. If this is occurring you will see a bright, diffused ghost image of the scene superimposed on the scene itself. Interference filters usually have significant reflections since the thin layers absorb little light, and where the filter is 50% transmitting at a passband edge it is also 50% reflecting. We chose not to tilt the detector, since that was mechanically difficult, but instead tilted the filter 2 degrees so the CCD would only see a black diffuse area reflected by the filter. As added insurance we painted the ceramic substrate of our CCDs with flat black paint and a very fine brush. We also tipped the polarization scrambler 2 degrees in the other direction, but we still ended up with a stray light path where reflections off the filter would reflect from the scrambler and produce a focused ghost 8 degrees off axis. Fortunately, these ghosts were less than 0.1% in brightness and do not integrate up to a large problem since they are of small angular extent. It would have been analytically easier to correct had the ghost been in closer alignment to the scene, though. It is impossible to correct ghosts that result from bright areas not in the image.

3.8 Reducing Filter and Scrambler Ghosts:

A ghost can occur when plane parallel filters are used since a reflection off the back side of the filter glass can be returned from the front side of the filter into near-alignment with the original light rays. To avoid this, we specified that the filters be made plane parallel to high accuracy, which would have confined a ghost to within a few pixels of the source target. However, to save money in fabricating 8 custom filters, the filters were made by sandwiching several pieces of glass, high and low pass filters, and blocking filters, and the sandwich was post-polished parallel. The unfortunate mistake we made is that the internal layers have significant reflectivity, and are not necessarily parallel to each other or the external surfaces. These buried layers create multiple ghosts. Figure 13 illustrates the ghosts seen when the instrument views a bright star-like source, for each band. In the figure, a 2% reflectance would be solid white. As you can see, we ended up with prominent, large ghosts. Each thumbnail image is 200x200 pixels in extent, to give the reader a sense of scale. For Unit Two (illustrated) the ghosts can extend up to 97 pixels from the source location. If a bright cloud in the scene under examination from orbit is small, this problem will be of little effect. For a large cloud bank or shoreline, however, it integrates up to a magnitude equal to 2.5% of the cloud band brightness – quite significant when viewing a dark ocean near the cloud bank. Bands 7 and 8 are particularly important since the ocean is quite dark at these wavelengths and they will be most corrupted by stray light from nearby jungle foliage.

Figure 13: Filter ghosts for Hawkeye Unit Two, for each Band



Typically in processing ocean color images one applies a cloud mask, to delete an area of the scene near a cloud from contributing to the data products. In this case that can be a significant portion of the image. So, we have been examining a computational approach to reduce this corruption. A computer program was written that can calculate a two-dimensional model of the scattering from our stray light measurements, and apply it to an arbitrary scene. First, an “impulse function” is calculated for each band, which is the amount of scattered light in neighboring pixels scaled by the total energy in the core of the image. Typical values are quite small, around 1 part in a million, so floating point math must be used. The core of the impulse function is zeroed out so only the stray light and ghosts remain. Then, the starting scene image is input and, for each pixel, the impulse function is overlaid on it and used to calculate what the impulse function contribution to nearby pixels would have been. These impulse function contributions are summed in a separate image file and saved. For a 200x200 pixel impulse function file and a 600x600 pixel image the algorithm takes about three minutes to compute. The astute reader will note that there is an inherent error in our thinking here, since, the original scene image is already corrupted, and yet we use it to calculate the corruption. This works since the stray light is only on the order of a few percent, and a 2% error times a 2% scattering level is insignificant.

This technique can be used to correct the stray light. In Figure 14-A I show a picture of the moon captured with the engineering model Band 6, with the contrast pushed hard. Figure Fourteen -B shows the impulse function for this band. Figure Fourteen -C shows the derived stray light estimate, and Figure Fourteen -D shows the original image with the estimate subtracted from it. All are displayed at the same contrast. This fairly simple technique reduces the stray light by a factor of 5 to 10 times.

Figure 14-A: EM - Lunar image with contrast pushed to show ghosts




Figure 14 -B: Impulse Function





Figure 14 -C: Derived Ghosts from Impulse Function



Figure 14 -D: Lunar Image Minus Derived Ghosts.



The better solution here would be to spend more money and procure filters that are plane parallel, with no buried layers. In retrospect we should have done that for Bands 7 and 8 at a minimum. However, we have observed in our test scenes of the ocean and lakes around Santa Barbara that the stray light is invisible due to the natural variations in the scene being much greater in magnitude.

4) FIELD TESTING OF SENSOR

Images from Field Tests: due to the small size of our instrument it is easy to mount it on a rotary platform, carry it to various locations, and capture multi-band images: Figure 15 is a natural color view of the city of Santa Barbara using Bands 6, 5, and 4. Figure 16 is a view toward the ocean and islands from Goleta. This image used Band 8 for the red channel, Band 7 for the green and band 6 for the blue. Figure 17 shows a closeup – it is impossible to show full resolution pictures of this size in a paper. Figure 18 is a view of the mountains behind Santa Barbara, and Figure 19 some detail from this scene. Except for dark subtraction and the application of flat field images to remove pixel-to-pixel variations (and some dust streaks) no other processing was done.

Figure 15: Santa Barbara as seen by Hawkeye – RGB = Bands 6, 5, and 4



Figure 16: Goleta and the Channel Islands, False Color – RGB = Bands 8, 7, and 6



Figure 17: Beach Area Detail



Figure 18: Mountain Scene – RGB = Bands 7, 6, and 5



Figure 19: Mountain Scene Detail



5) OBTAINING DETAILED DESIGN INFORMATION

For more design information: per the Moore Foundation contract, all design information, software listings, mechanical drawings, schematics and such are available to qualified researchers through John Morrison at the University of North Carolina – Wilmington. Most of the software tools necessary to best view these materials are available from third parties. Image analysis and processing was done using SBIG's CCDOPS program, which now has very limited availability. Cyanogen in Ottawa, Canada offers an archived version of CCDOPS, as well as a similar program, Maxim DL.