JOHNS HOPKINS UNIVERSITY **Applied Physics Laboratory**
Laurel MD 20723-6099

A1F(2)02-U-067
October 29, 2002
ICY0243X

To: S. L. Murchie

From: L. M. Howser

Subject: Analysis of Noise in the CONTOUR Spectrograph Images

Summary

Images from the CONTOUR Remote Imager and Spectrograph (CRISP) calibration tests showed a higher-than-expected level of noise. Analysis was performed on images from the cooler cycling tests to determine the statistics of the noise and to help ascertain the source of the noise. This analysis included searching the header files of the images for unique values, and determining the frequencies present in the noisy images. Spectrograph data from the spacecraft thermal vacuum tests were also included in this analysis. The calibration images showed two different noise patterns; the spacecraft noise patterns did not match either of the calibration noise patterns. Because the CRISP instrument was not in the final flight configuration during the calibration tests, the noise measured during the thermal vacuum tests is expected to be more representative of the noise in flight conditions.

Introduction

Three sets of data were used in this analysis. Two of the datasets were collected during the cooler cycling calibration tests at APL. The third dataset was collected during the spacecraft thermal vacuum testing.

The cooler cycling tests were run to determine whether vibration from the spectrograph cooler caused any noticeable vibration in a point source image. Because of errors in test setup, two runs were made for this test. There was little difference in the test method between the two runs. However, images from the two test cases showed significant differences in the level of noise. These two cases were studied to find differences in the testing that may have changed the noise environment of the sensor.

During the calibration tests at APL, the CRISP instrument was not in the final configuration. A lab data processing unit (DPU), in a VME chassis, was used with the instrument because the flight DPU had already been integrated with the spacecraft. Therefore, flight shielding and grounding were not present during the calibration tests. The test equipment (motors, pumps, etc.) also generated noise that could affect the data collected.

As a comparison to the calibration test data, images from the spacecraft thermal vacuum testing were also analyzed. During these tests, the spacecraft was in the final flight configuration and the thermal chamber was cooled.

The noise in images from these three datasets was compared in both the shorter and longer wavelength bands of the instrument. The spatial and temporal standard deviation of the images and the FFT were used in these comparisons.

Calibration Test Noise Images

The goal of the cooler cycling tests was to determine whether the test equipment motion stage or the spectrograph cooler caused vibration in the spectrograph that could cause the image of a point source to move in the FPA image. During the first test run, the light source was not aligned with the imager. Data were collected with the motion stage and cooler ON and OFF, but there was no point source in the images. The test sequence was repeated with a point source in the field of view. The images used in this analysis were from a series of 'dark' images collected prior to the point source tests. No light source was on when these images were collected.

A columnar pattern of noise was present in the first set of data. Figure 1 shows an example of this noise. In these data, the column mean has been removed, thus removing spectral variation. The noise appears to be about ±4 counts. The phasing of the noise changes at a period of about 40 pixels. The trace at the bottom of the figure suggests that the noise is from a source with frequencies higher than the pixel clock rate.

It should be noted that the noise in the four quadrants is synchronized. Figure 2 shows an example from frame 1 of this sequence. The top plot contains traces from the left and right halves of row 52; the bottom plot has the corresponding traces from row 180. These two rows are 128 rows apart, so these four regions are read from the array at the same time. This figure shows that the phasing of the odd/even noise, and the 'beat' location is the same in all 4 traces, implying that the noise is added during the read-out of the data from the FPA.

One goal of this analysis was to determine the frequencies of the noise, thus aiding in finding the noise sources. Knowledge of the timing of the pixel clock was needed to accurately duplicate the noise pattern. The CRISP image is divided into 4 quadrants; data from the quadrants are read in parallel. Thus, pixels (1,1), (1,129), (129,1), and (129,129) are read out at the same time. The data are read out across a row of the quadrant at a pixel clock rate of 3.5 μsec / pixel. The read-out is then paused for reset, before the next line is read. Each line requires a total of 1.344 msec before the next line is read. After the line read-out, there is an overhead time on the order of a few microseconds.

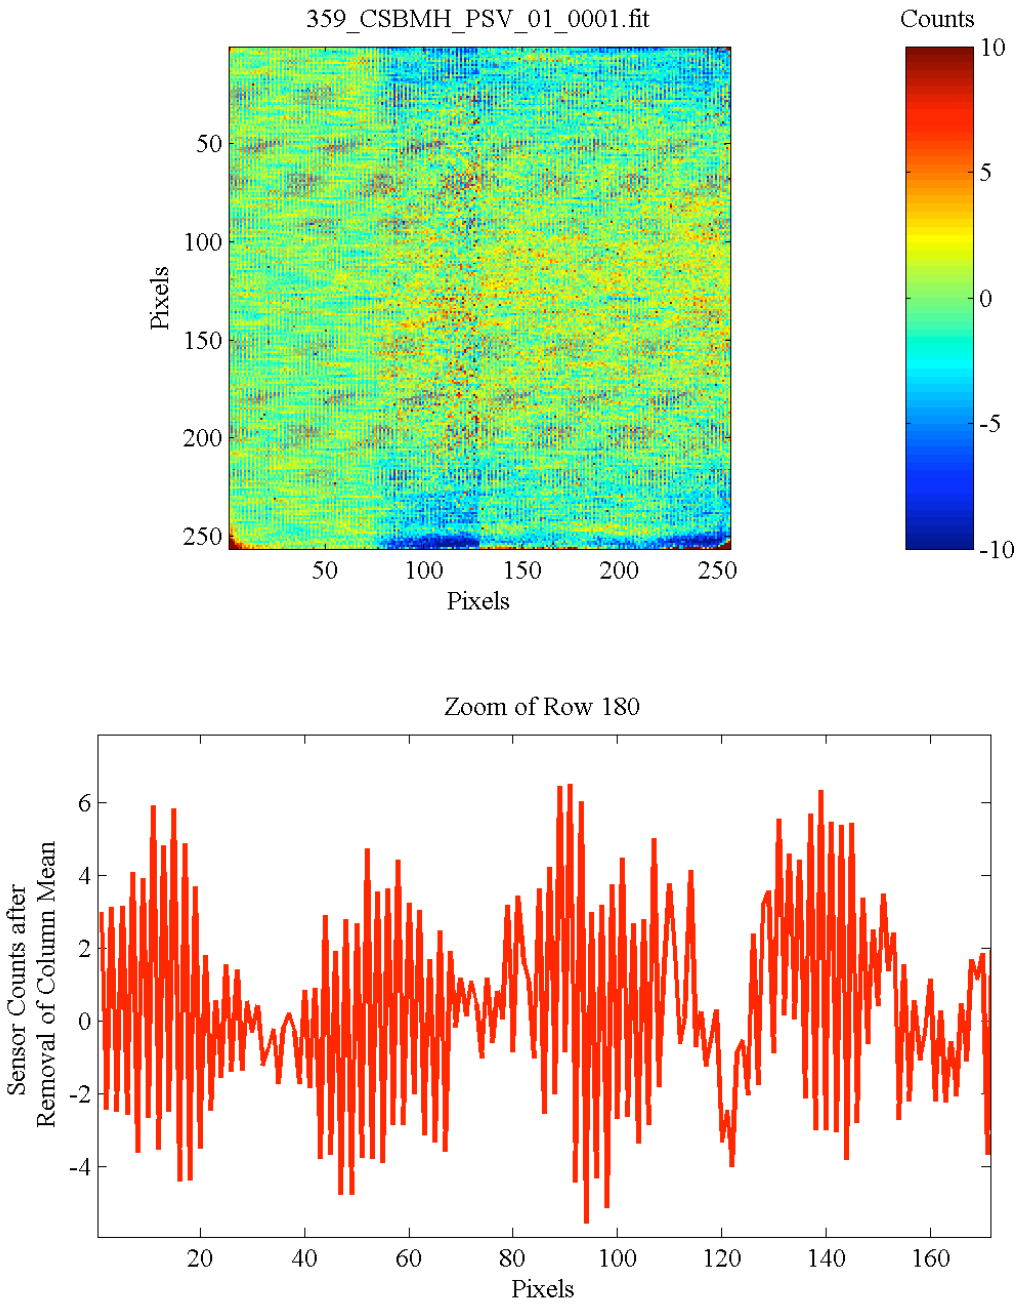

Figure 1: Example of column noise from run 1. The top half of the figure shows an image after removal of the column mean (removing spectral variations). The bottom half has a trace of row 180 from the figure (only pixels 1-170 were plotted for viewing purposes). The trace shows a column-to-column noise that appears to be at a higher frequency than the pixel clock (i.e., aliased from a higher frequency). The period between cycles is about 40 pixels.

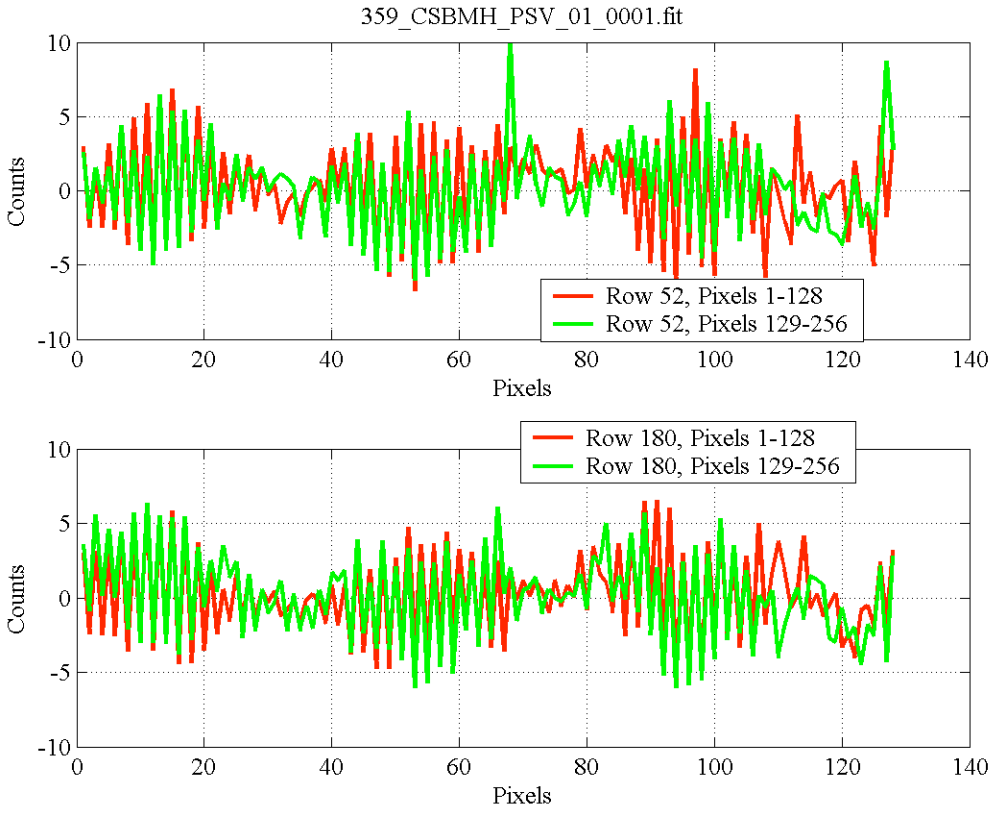

Figure 2: Horizontal traces from four regions of the array after removal of the column mean. The top two traces are from the left and right halves of row 52; the bottom traces are from the left and right halves of row 180 (128 rows below row 52). The timing of the FPA read-out is such that these four lines are read from the array simultaneously. Since the noise on the four traces is synchronized, it indicates that the noise occurs during the read-out, not during the image integration.

To duplicate the CRISP image noise, a sinusoidal wave image was generated in Matlab. The frequency and amplitude of the sinusoid were chosen so that the synthetic noise pattern would match the CRISP noise image. Figure 3 summarizes the results. The top 2 images are the synthetic sinusoidal image on the left and the CRISP noise image on the right. Both of these images have the column noise, along with a diagonal pattern from the beat frequencies of the noise. Horizontal traces of these images are plotted in the middle of the figure. These traces show the similarity in the beat frequency and the amplitude of the signal. Vertical traces of the images are plotted in the bottom of the figure. The frequency selected for the synthetic image was almost 425 kHz. This frequency is just under the 3 times pixel clock Nyquist frequency of 142.8 kHz. Note that very small changes in the frequency caused large changes in the resulting image. Also, any uncertainty in the knowledge of the pixel clock timing (including the

read-out overhead) would also have large affects in the noise image. Because the frequency is above Nyquist, any source at another 'Nyquist multiple' could also be a source of the noise.

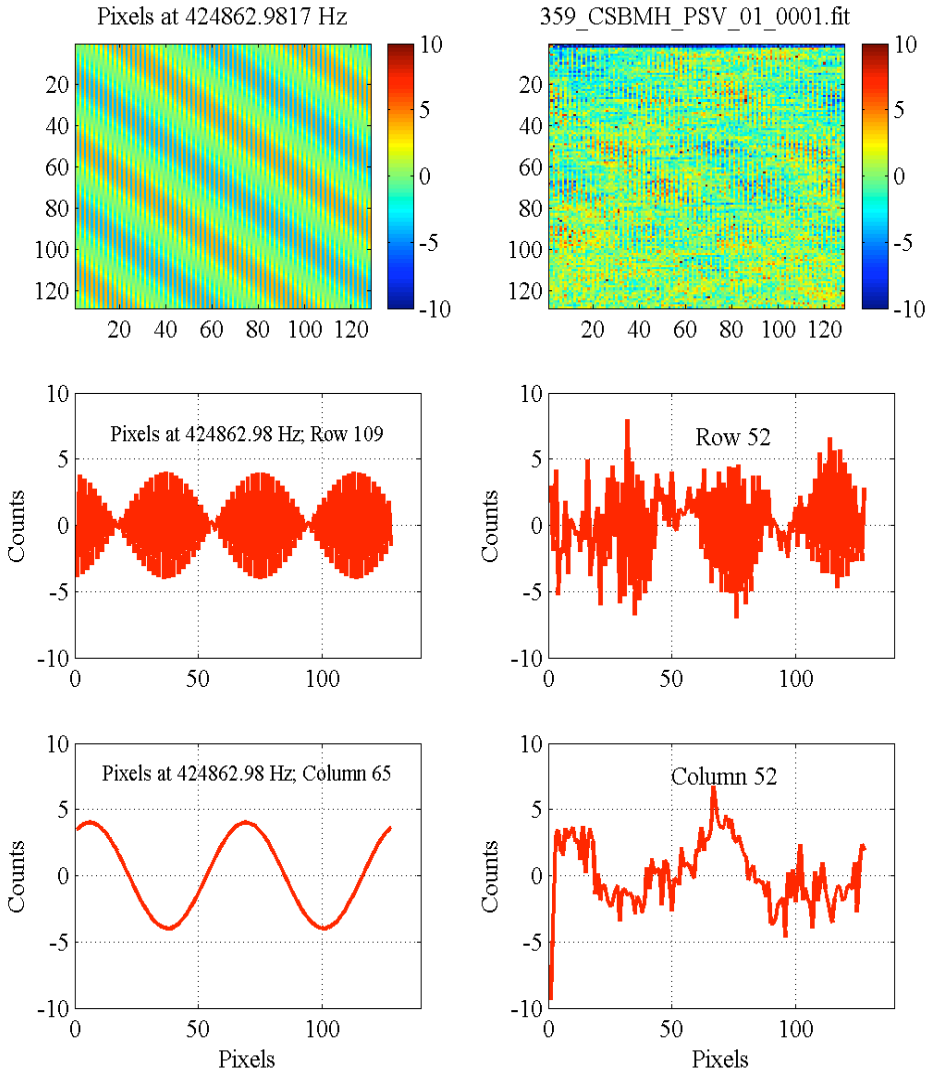

Figure 3: Comparison of a synthetic sinusoidal image with the CRISP spectrograph noise image. The sinusoidal image (at frequency of ≈425 kHz), matches the amplitude and beat frequency of the CRISP image. The horizontal (middle) and vertical (bottom) traces of the sinusoidal image also match the CRISP traces.

Figure 4 shows an image frame from the second run of the cooler tests. The column means have been removed from the image. The horizontal trace in the bottom of the figure shows noise spikes with amplitudes of ±20 counts. The odd/even column noise is not apparent in this figure. Similar to the data from the first run, the noise spikes in this image are also synchronized over the four quadrants.

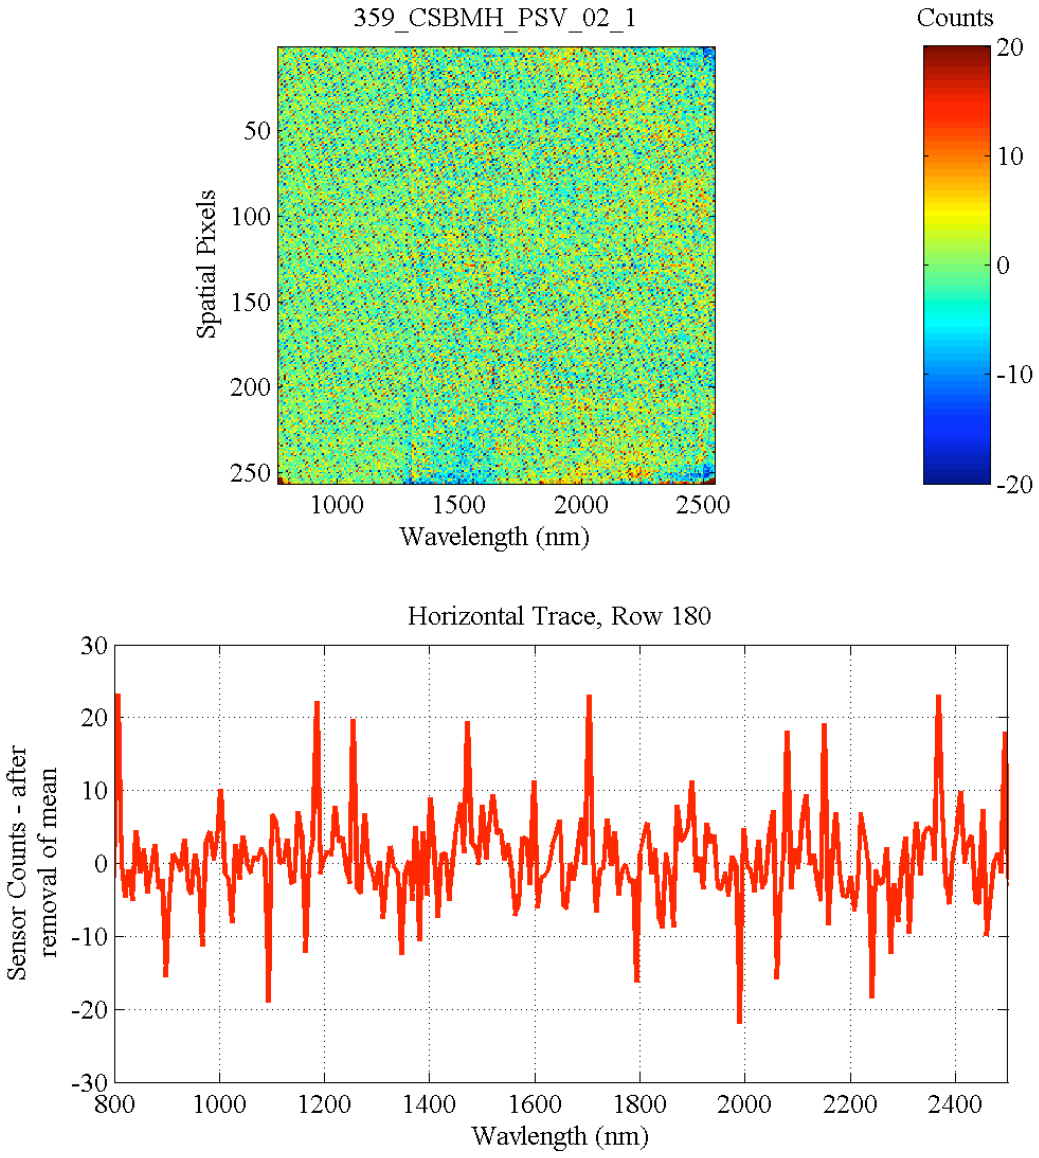

Figure 4: Image from the second run of the cooler test. This image and the corresponding horizontal trace, show higher amplitude noise than the data from the first cooler test.

An FFT of the images from the second run of the calibration tests showed that there was no single dominant frequency. Therefore, instead of determining the noise frequencies with sinusoidal images, the FFT results were used directly. Figure 5 shows the results of this analysis. Prior to computing the FFT, the average of all frames in the run was computed. This average image was subtracted from the first frame. Thus, the spatial noise was removed prior to the FFT. The FFT was the computed on a single line of the image. As a comparison, the FFT results from Run 1 are included in the plot. The figure shows that the odd/even noise, aliased down to about 140 kHz, is present in both datasets. However, Run 2 has many frequencies with higher magnitudes. In addition, both runs show some energy at 5 kHz.

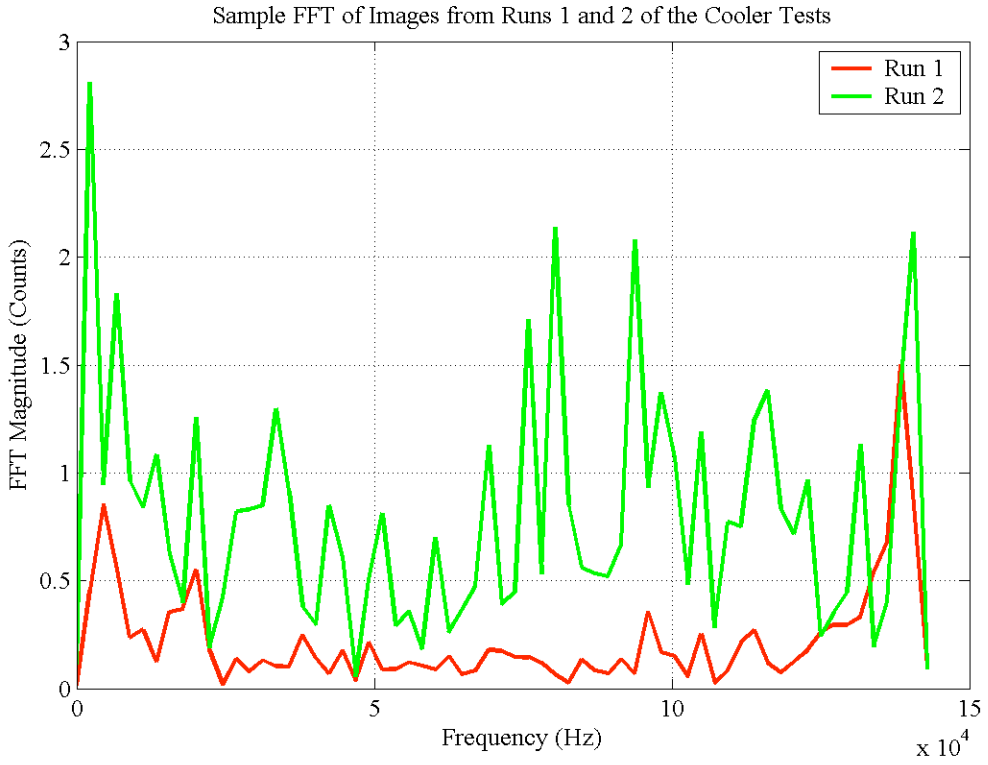

Figure 5: FFT of images from Runs 1 and 2 of the calibration cooler tests. These plots show the higher noise in the images from run 2. The odd/even column noise is shown at about 140 kHz.

Calibration Test Conditions

The previous analysis showed significant differences in the noise between runs 1 and 2 of the cooler cycling calibration tests. Analysis of the header information was conducted to determine the differences in the test conditions between the two runs. The following table summarizes the results. Seventy-one variables from the image headers are listed in the table. These values were read from each of the images in the two runs; averages are listed in the table.

Table 1 Summary of Header Data

Variable	Run 1	Run 2	Comments
Mirror Motor Current	≈2100	≈2100	
Mirror Heater	0	0	
Mirror Position	**48982 (-93.1°)**	**58003 (-90°)**	***
Mirror Motor Heater Current	≈65000	**≈65000; with spikes down to 0**	***
Mirror Mode	2	2	
Cool Current	**500**	**5; with spikes to 65000**	***
Cool Converter Current	≈600	≈600	
Cool voltage	**≈2000**	**≈5; with spikes to 65000**	***
Cooler power	1	1	
Cooler Temp	1	1	
FW 15v Current	≈200	≈200	
FW 15v voltage	≈6100	≈6100	
FW converter Current	≈400	≈400	
FW Current	≈100	≈100	Both runs had spikes up to 65000
FW motor power level	4	4	
FW phase	**0**	**1**	***
FW position	≈58400	≈58420	
Fw motor – motor current enabled	0	0	
FW motor board primary power	1	1	
Fw pri current	≈450	≈450	
Fw resolver	1	1	
Filter	6	6	
Frame rate	5	5	
Atten pos 1	0	0	
Atten pos 2	55	55	
Auto flush	1	1	
Bulk heater	0	0	
Bulk heater current	170	170; large spikes to 300	
CA mode	0	0	
CA state	0	0	
DPU Current	5800	5800	
DPU voltage	1500	1550	

Table 1 (cont'd.)

HOP1 heater1	40	40	
HOP1_1	0	0	
HOP1 heater2	40	40	
HOP1_2	0	0	
HOP2 heater 1	40	40	
HOP2_1	0	0	
HOP2 heater2	40	40	
HOP2_2	0	0	
Diaphragm heater	0	0	
Diaphragm heater current	≈150	≈150	
Downlink	1	1	
Image converter	≈770	≈770	
Imager current	≈1150	≈1150	
Imager voltage	≈6000	≈6000	
Image power	1	1	
Sp power	1	1	
Sp primary power	1	1	
Spec converter current	≈720	≈720; with spikes up to 1000	
Spec current	≈1200	≈1200; with spikes up to 1400	
Spec primary current	≈1400	≈1400	
Spec voltage	**≈5000**	**≈4000 and noisy**	***
Star1 heater	0	0	
Star1 heater current	≈170	≈170	
Star 2 heater	0	0	
Star 2 heater current	≈140	≈140	
Tracking Mirror Temp (°C)	-33	-26	
Upper Housing Temp (°C)	-33	-26	
Lower Housing Temp	-32	-24	
Scan Motor Temp	-32	-23	
Table Temp	-30.4	-30.5	
Liner 3 o'clock	-35.7	-35.6	
Liner 6 o'clock	-36.2	-36	
Pitch stage	-30.9	-30	
Chamber wall	22	23	
Cooler Return	-31.3	-31.1	
Cooler supply	-33.8	-33.6	
Target chamber vac (Torr)	6.2e-8	9.7e-8	
Collimator vacuum	1.25e-7	1.65e-7	
Monochromator vac.	2.3e-6	2.18e-6	

A number of issues are evident in the table. Of primary interest are the variables that differ between the runs. These variables are in BOLD in the table. They are:

1) the mirror position – changing from –93.1° (run 1) to –90° (run 2),
2) the cooler current – changing from an average of about 500 to about 5 counts,
3) the cooler voltage – changing from an average of 2000 to an average of 5,
4) the filter wheel phase – changing from 0 to 1, and
5) the spectrometer voltage – changing from 5000 to 4000.

In addition to these issues, many of the variables were noisier in run 2 than in run 1. Also, a number of variables showed large jumps from near 0 to over 65,000. It appears that these variables were incorrectly converted before being placed in the image header structure. A 2's complement conversion to positive and negative values would probably fix these variables. A final observation is that many of the temperatures reported from run 1 are a few degrees cooler than reported from run 2.

Spacecraft Thermal Vacuum Data

Similar image sequences were collected during the spacecraft thermal vacuum tests this past spring. On March 22nd, dark images (no light source) were collected when the thermal chamber was cooled and while it was allowed to warm up. Because of test issues, this warm-up sequence of CRISP images was only collected for about a half hour. Images were collected at 1, 2, 3, 4, and 5 Hz frame rates. Since all of the calibration images were collected at 5 Hz, we'll concentrate on the 5 Hz data from the spacecraft data.

Three examples of the spacecraft images are shown in Figure 6. During this test sequence, three sets of data were collected as the sensor warmed up. The spectrograph housing temperatures for the three examples were: -26.26°C, -22.75°C, and -19.25°C. The plots on the right of the figure are horizontal traces from a single row of the images on the left. The temperature-induced increase in signal in the longer wavelengths is evident in the three images. The average signal in the middle temperature image most closely matched the signal in the second calibration run.

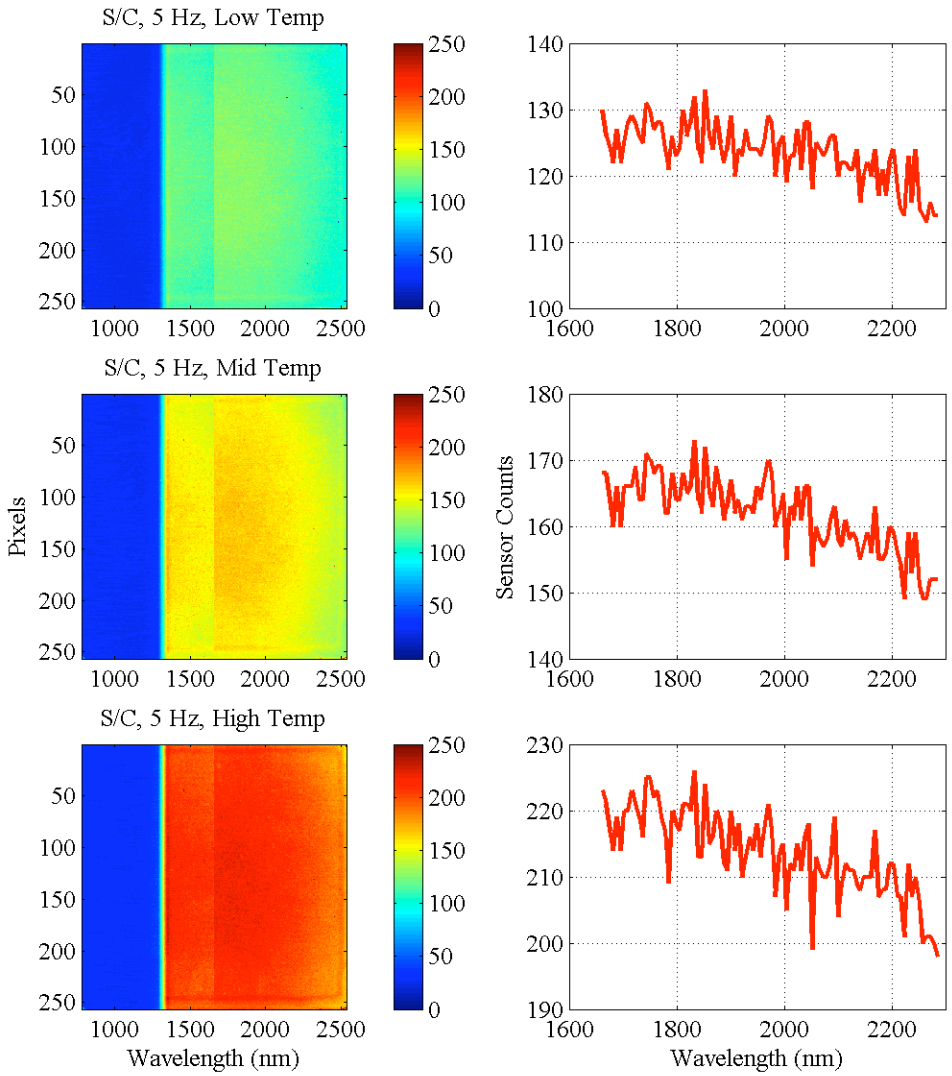

Figure 6: Example spectrograph images from the spacecraft tests. These images show an increase in signal as the sensor warmed up. The horizontal traces on the right show a corresponding increase in noise along with the increased signal.

A comparison of the noise from the three datasets is shown in Figure 7. The images on the left of the figure have had the column mean removed (thus removing the spectral signal). After removing the column mean, it was noticed that the image was also non-uniform in the spatial dimension of the longer wavelengths. For the long wavelengths, the image was brighter in the center than at the edges. Thus, the mean of each row was also removed from the longer wavelength data. In the images in this figure, the low frequency trends were removed, only the higher frequency noise remains.

The plots on the right of the figure show a horizontal trace from a single row of the image. These traces show that the 2nd calibration run had the highest noise. The noise in the images from the spacecraft thermal vacuum tests was NOT synchronized between the four image quadrants.

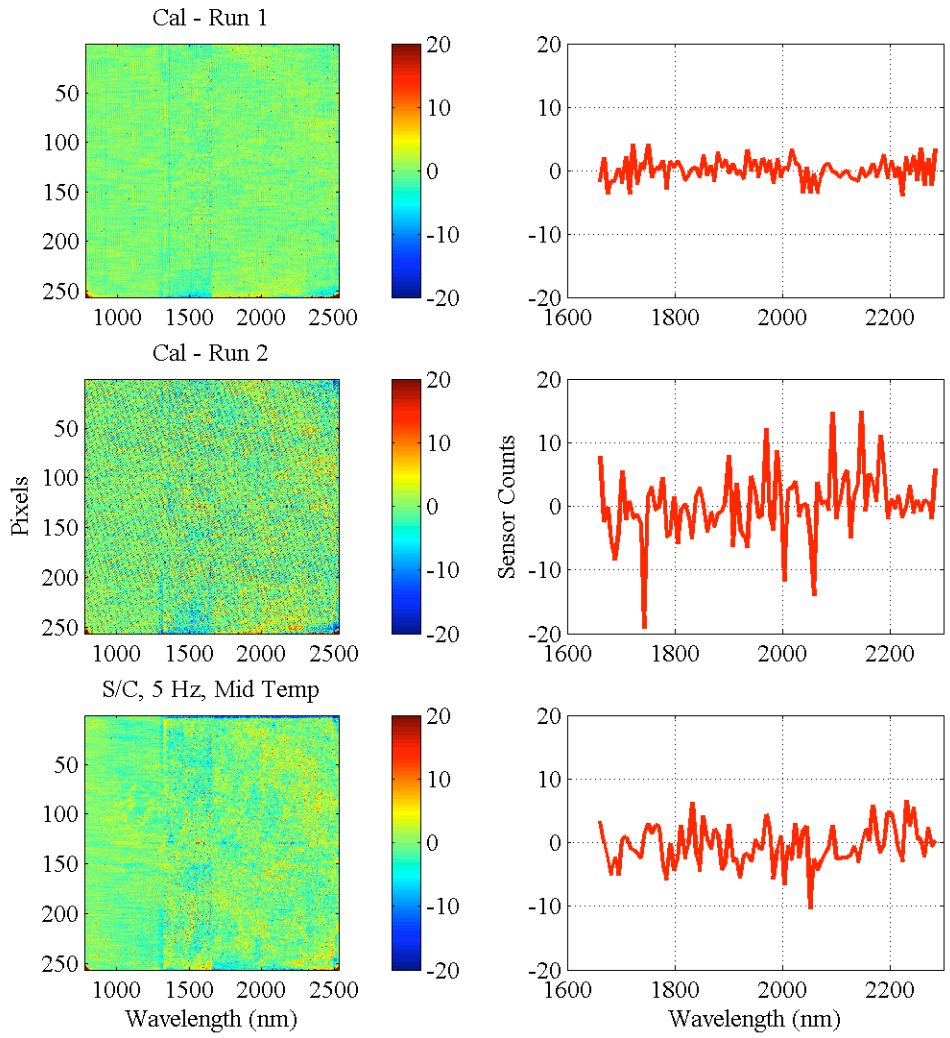

Figure 7: Comparison noise images from the two cooler cycling calibration runs and the spacecraft tests. The images on the left have had both the column (spectral) and row (spatial non-uniformity) means removed. Thus, only the higher frequency noise remains. The plots on the right show horizontal traces from a portion of a single row of each image. The lower noise level in the first calibration run is evident.

Figure 8 shows a comparison of the standard deviation of the images in the three tests. (The bottom of the figure is a zoomed version of the top.) In this figure, the standard deviation of the image is plotted versus the average scene intensity. The standard deviation was computed after the row and column means were removed. Both variables were computed in the short wavelength and long wavelength bands. The spacecraft tests had data at three frame rates and three sensor temperatures. These data show a curve with increasing noise as the scene intensity increases. In contrast, images from the first calibration run had higher noise in the shorter wavelength (lower intensity) region. Images from the second calibration run had significantly higher noise than the spacecraft tests, with both wavebands showing similar noise levels. Note that the spacecraft noise in these plots was dominated by spatial (gain) nonuniformities in the image.

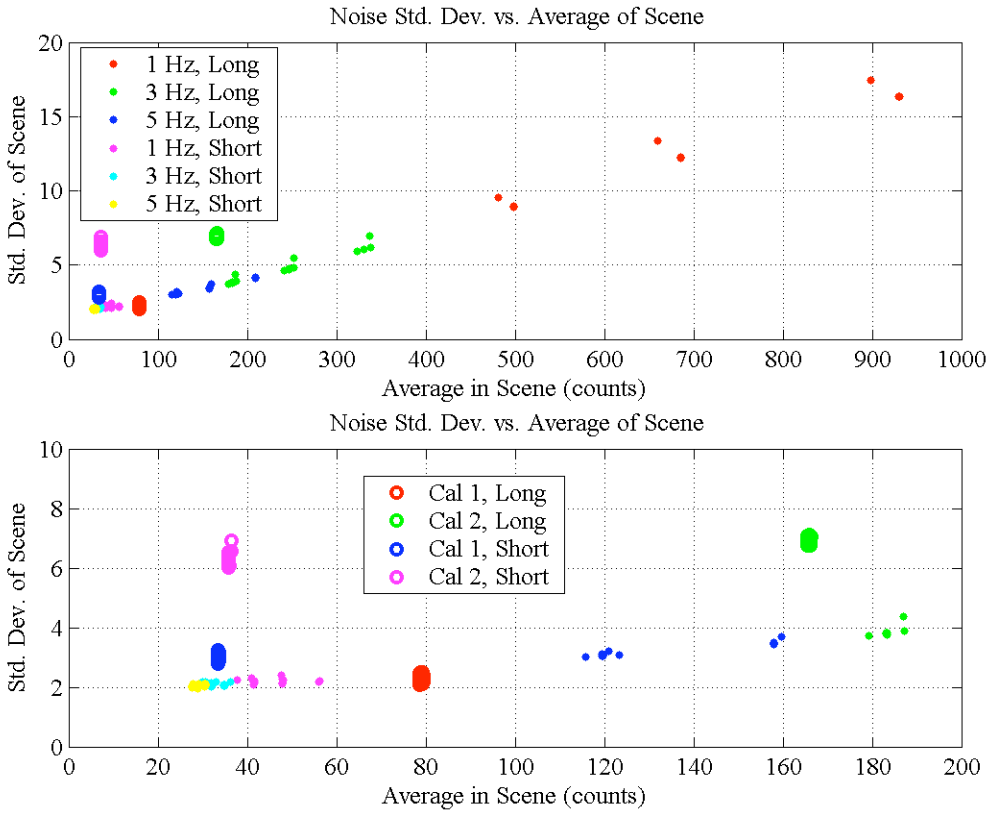

Figure 8: Noise standard deviation versus average scene. These plots show the increase in noise as the scene intensity increases. The increase in noise with instrument temperature is visible in the spacecraft data. Data from the second calibration run are significantly noisier than the other data. (The bottom plot is a close-up of a region of the top plot.) Since the standard deviation was computed across a region of each image, this noise is probably dominated by spatial nonuniformity at the higher scene intensities.

Temporal noise in the image sequences was calculated by computing the standard deviation of the signal from each pixel (in time). This calculation generated a 256 by 256 image of standard deviations. The temporal noise was estimated using the median of this temporal noise image. The results of this calculation are shown in Figure 9. The temporal noise in the spacecraft images is significantly lower than the noise in the calibration data. This figure confirms that the spacecraft noise in Figure 8 is dominated by spatial (gain) nonuniformity.

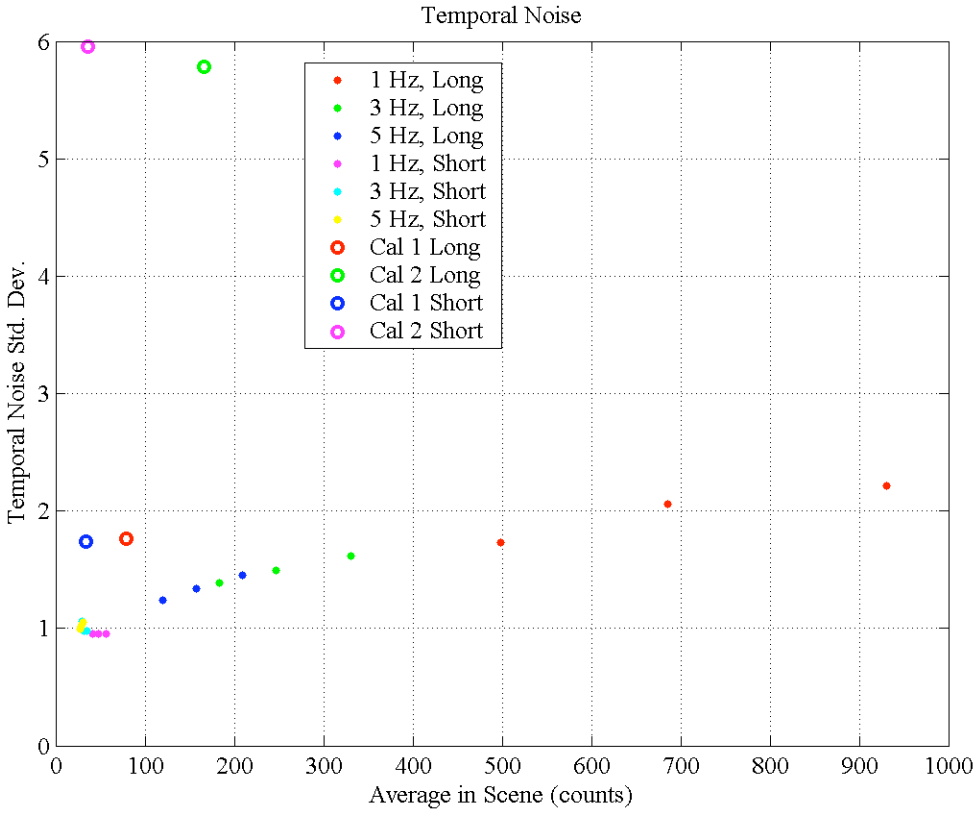

Figure 9: Temporal noise versus average in scene. This plot shows that the temporal noise of the spacecraft tests (1, 3, 5 Hz data) is significantly lower than the noise from the calibration tests. The spacecraft temporal noise is linear with the square root of the average scene intensity.

Figure 10 compares the FFTs of the 5 Hz spacecraft data with the calibration data. As in Figure 5, these FFTs were computed after removing the spatial non-uniformity data. As seen previously, the second calibration run has more energy than the spacecraft data. The first calibration run and the spacecraft data are similar except the aliased data near 140 kHz are not present in the spacecraft data. All three datasets appear to have some energy near 5 kHz.

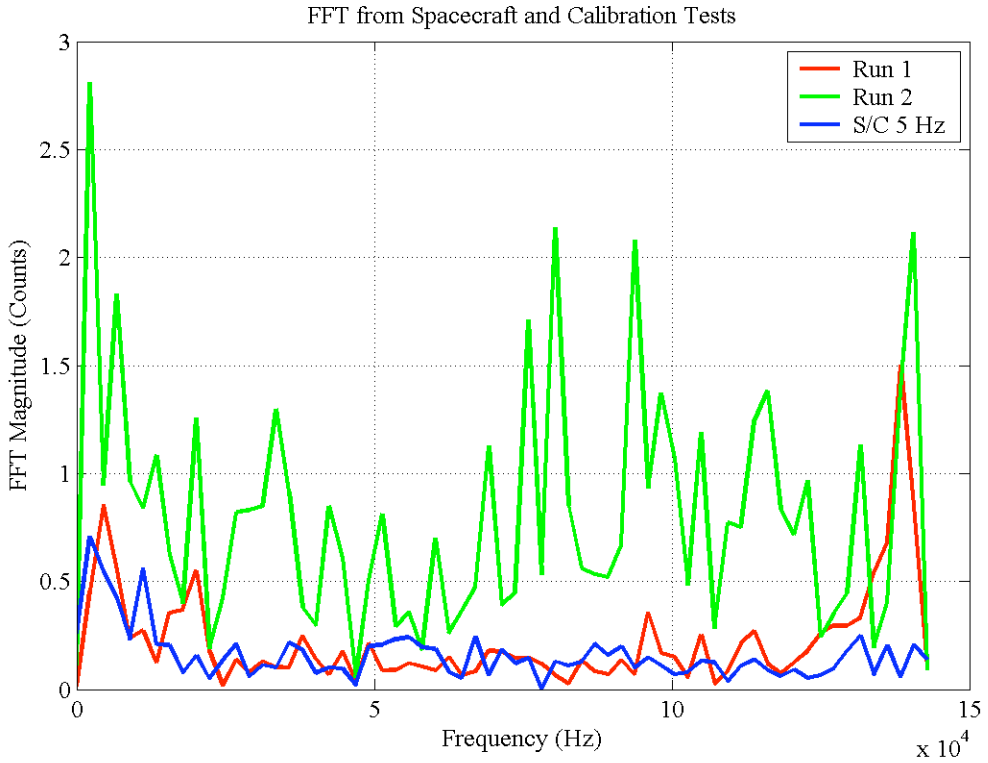

Figure 10: FFTs of the data from the spacecraft tests and the calibration tests. An FFT from the middle temperature, 5 Hz spacecraft images is compared to the calibration data.

Conclusions

Data from the two cooler cycle calibration test runs were analyzed to investigate the noise sources in the CONTOUR spectrograph. The first cooler run had odd/even column noise at a frequency probably aliased down to 140 kHz from higher frequencies. Analysis showed this noise may be at a frequency of about 425 kHz. In addition to the column noise, the images from the second run had significantly higher noise at multiple frequencies. Five variables in the header file showed some difference between the two runs. These variables may suggest a source of the image noise. Because the instrument was not in the flight configuration during the calibration tests, the calibration data are not representative of the noise expected in flight.

The spacecraft thermal vacuum data had different noise than the calibration runs. The odd/even column noise was not visible in the spacecraft noise. The first calibration run and the spacecraft data showed a minimum noise standard deviation of about 2 counts for the lowest light levels. Higher input signal levels had high noise, caused by spatial non-uniformity. Temporal noise of the spacecraft images is shown to be linear with the square root of the scene intensity. Temporal noise is significantly higher in the calibration images than in the spacecraft images.

L. M. Howser

Distribution

S. A. Gearhart
B. L. Gotwols
L. M. Howser
D. C. Humm
N. R. Izenberg
M. J. Mayr
R.W. McEntire
S. L. Murchie
E. W. Rogala
W. J. Tropf
J. W. Warren
A1F Files
Archives