STARDUST NAVCAM DATA

I. Background

The STARDUST navigation camera admirably served its primary purpose of allowing optical navigation for the encounter with P/Wild 2. The comet was detected on the first attempt, made six weeks before the encounter. During the encounter, the camera was used for scientific study of the comet's nucleus, and significant new knowledge was gained on the morphology of the comet nucleus and the bearing strength of its surface. It has proven much more difficult to use the data for absolute photometry. Treatment of these data up to mid-June is considered in this summary.

Extensive laboratory tests were run at JPL before the camera was delivered to Lockheed Martin for installation on the spacecraft. The absolute sensitivity of the camera was determined using a collimated beam supplied by an integrating sphere having a diameter of 30-inches, with the beam output determined by an NBS calibrated radiometer. The beam output was uniform to better than 1% over the entire part of the beam incident upon the camera. The camera and its attached scan mirror were inside a vacuum chamber, and the resulting response of the CCD detector was measured at temperatures of –30, -40, and –50 degrees centigrade.

The periscope, which protected the camera from the cometary dust during encounter, was delivered six weeks late, so the two systems were never tested together on the ground. The reflectivity of the periscope mirrors was measured, and the periscope was installed on the spacecraft at the last minute. It was not perfectly aligned, however, and this resulted in double images when part of the light came through the periscope and part did not. The two images are separated by about 17 pixels in the focal plane and occur whenever the scan mirror feeding the camera was at angles between 6 and 17 degrees. This cost us 13 of the 72 images for photometric purposes. However, a check after the encounter showed that the periscope had done its job very well. Its mirrors were sandblasted, while the camera and scan mirror were as good as ever.

The first images were attempted 2½ months after launch. After taking two images, the spacecraft safed itself. The two images appeared very peculiar, showing nothing that looked like a star field. The next attempt was made 10½ months after launch, and it was immediately obvious that the camera was heavily contaminated by an unknown substance or substances. Only one star was readily visible, and it appeared 100 times (five magnitudes) fainter than it should have. Turning on the camera heaters brought it up above zero centigrade and resulted in some improvement. Finally, turning the Sun on the radiative cooler brought the detector up to 30 centigrade, and the images improved enormously. When another series of images (of the Moon) was attempted two months later, some of the contamination had returned. Thereafter, whenever a critical set of images was to be attempted, decontamination was performed. All of this meant that any thought of using the original laboratory absolute photometric calibration was hopeless. Calibration, as described in a later section, was performed in flight.

II. Standard Data Processing

Most of the image processing has been done using Excel, after using ProView to read the pds files from the STARDUST data computer. Unfortunately, although Excel has some 64,000 rows, it has only 256 columns. A ProView program was written to split the images into four parts, called A, B, C, and D, each 255x1024 in size. These could then be read and manipulated in Excel. The 256th row in part D of each image was a fat pixel and useless. In addition, a final free column in each piece was useful for such manipulations as summing rows. The first column in part A was also a fat pixel and was ignored and not used in calculations. After all work on each part of an image was completed, we were able to reassemble the parts into a single processed image that could be displayed as an image. So far this has been done for only one image. In a casual viewing, the processed image appears little different from the original compressed image, but these linear images can now be studied quantitatively to derive properties of the Wild 2 nucleus.

Three standard procedures were used on all data. The data were recorded compressed by a hard-wired square root compressor, and images were therefore non-linear in their exhibited intensity. Compressed dn were changed to uncompressed using a lookup table. Each compressed level was equated to the middle of the corresponding uncompressed range. A copy of the lookup table is attached as Appendix I.

The background of this or any imaging camera and CCD detector changes with temperature and to a lesser extent with aging. A positive electronic bias level was preset to avoid any possibility of the signal going to negative values. That bias level changes with temperature, perhaps counter-intuitively increasing as everything gets colder. That bias level can be determined by recording zero exposure frames, but it can also be determined from the so-called BLS (Base Line Stabilization) pixels. A raw compressed image from our 1024x1024 pixel CCD actually has 1048 columns, with an additional 13 pixels before each line and 13 at the end of each line. The first two columns are a sync word followed by two columns for the line count. Then come the 8 BLS pixels. Finally there is a so-called fat pixel, which is the first column of the image. In each case, these BLS pixels are equal to the bias level. Over time, if there is a change in the efficiency with which the CCD transfers charge from one column to the next, it will show up in a difference between these leading BLS pixels and the 12 BLS pixels which follow the image columns, or if the degradation is serious, even from one column to the next of the exhibited value of the BLS pixels. The “bottom line” is that this bias must be subtracted from each pixel in the image, since it is an electronic artifact and not caused by incident photons.

In any system utilizing transmission optics, there will be a small variation in the total transmission from the center to the edges. This generally amounts to at most a few percent but can be corrected. The absolute pre-launch data provide a comparison between the central pixels and all other pixels. We divided all images by the ratio of the average of the 100 central pixels (10x10) to all the other pixels to remove this vignetting. The vignetting matrices are included in the email transmission as XratioX.xls, where X can be

A, B, C, or D for the four image pieces. (These are obviously Excel files.) The vignetting should remain the same over time except for possible irregular deposits of contamination. We chose to apply this correction to all images, assuming we had removed all or most of the contamination with heating. This treatment will also remove any new "hot" pixels caused by particle irradiation of the detector. We experienced several solar flares during the five year flight from Earth to Wild 2. In most cases the pixels return to their normal sensitivity after a few days, however.

The first and last columns of an image are so-called "fat" pixels. These must simply be ignored, having a much larger value than the light incident upon them would warrant. The useful image therefore contains only 1022 useful columns. The STARDUST camera has an angular resolution of 59 microrad/pixel (12"/pixel) and a focal length of 202 mm at an f–ratio of about f/3.5. Early in the flight the filter wheel failed, possibly due to a failed power supply. Fortunately it failed on the filter with the largest throughput, but the broad bandpass of that filter caused images taken through it to have significant chromatic aberration, which resulted in an image resolution of about 2.5 pixels at FWHM (full width at half maximum) when observing a point source such as a star. (The high resolution filter, intended to be used for near encounter imaging, would have resulted in resolution exceeding a half pixel.) Without any image processing, the 2.5 pixel resolution resulted in a best linear resolution at closest approach of about 20m/pixel.

III. Image Reduction & Wild 2 Properties.

There is a background in each image, even after the bias is removed. This can originate from any one or more of at least four sources. The comet has a coma of dust and gas that is very large when compared to the size of the nucleus and therefore quite uniform radially in our small field of view. (There are azimuthal variations, however, depending upon the amount of solar radiation incident upon the local surface and the distribution of dust sources.) There is scattered light from internal sources, largely caused by imperfect anti-reflection coatings on the lenses and by contamination of the optics as noted in the Background section. There is scattered light from external sources. The largest of these is reflection from the rear of the spacecraft, especially the sample return capsule, when the scan mirror angle is 160 degrees or more. At small mirror angles (less than 10 degrees), there is evidence of light scattered into the periscope, perhaps from the solar panels. The evidence of scattered light at either extreme of the scan mirror travel is the existence of a minimum in the signal away from the nucleus followed by an increase with increasing distance in the field of view from the nucleus. Finally there is the possibility of some biasing in decompression of the data, caused by using the midpoint of the uncompressed range.

In near-nucleus (middle) images, the level of the background was determined by finding the mean level of as many as the six columns (6144 pixels) farthest from the nucleus or as few as 10 rows and six columns (60 pixels), in the corner farthest from the nucleus, this at the times when the angular size of the nucleus is at its greatest (near closest approach). The mean background is usually between 5 and 10 dn in central parts of the field, but the largest images do have greater values, suggesting that internal scattered light may be making a contribution. The uniformity among many of the images taken at middle parts of the scan mirror travel seems to suggest coma as the source and to rule out scattered light as the principal contributor in these cases.

For the initial studies of nucleus size, mean albedo, and phase function, two quantities are needed, the size of the illuminated part of the nucleus and the total amount of light reflected or scattered from the nucleus. The latter, of course, includes the total dn from the actual area of the nucleus as well as light scattered by the nucleus to points outside the nucleus proper. This is measured as the total dn above the coma or background. If another source is contributing to this measurement, then the internal scattered light is overestimated, but there is little if any evidence of significant external contribution when the scan mirror angle is between 20 and 150 to 160 degrees.

The cross-sectional area of the illuminated part of the nucleus has been determined using the Excel countif function. The columns and rows limiting the greatest physical extent of the nucleus are determined, and the pixels are summed for all values greater than the background level between these greatest coordinates. The edge coordinates are determined by extending the nearly linear drop in dn in the outer part of the edge to the background level. The countif function then sums all pixels above the background in an area of 4ab, where 2a and 2b are the sides of the rectangle defined by the four edges. Most of the images are fairly elliptical in shape. The area of an ellipse with semi-major axes a and b is πab, a much better estimate of the cross-section of the nucleus. Therefore I multiply the rectangular area determined from the countif function by the ratio $\pi/4$ and use this as my area estimate. I realize that there are much more sophisticated ways to model the edge, but this is a fair approximation when time is limited. There are what look like nearly disconnected pieces of the nucleus that rotate into view as the spacecraft passes the nucleus. There are gaps that appear nearly black between these pieces and the nucleus, but in fact they are at a dn level much greater than the background level, so I have assumed that they should be included as part of the illuminated nucleus. I would be <u>very</u> happy to have someone improve on this admittedly crude approximation.

The total irradiance of the nucleus is computed by adding up the contributions of reflected light from all of the pixels on the visible portion of the nucleus. First, the pixel corrected from dn to dn/s by multiplying by 1000 over the exposure time in milliseconds. It is then corrected to a standard distance of 1000 km by multiplying by the square of the range at the time of the exposure divided by the square of 1000. The linear size of a pixel at a range r is just r in km times the 59 μrad/pixel angular resolution of the camera. The area of the illuminated nucleus in km^2 is just the number of pixels multiplied by the square of the linear resolution. Finally, the product of <u>mean</u> albedo times the phase function at a particular phase angle is given by the simple photometric equation given by Russell in 1916 (ApJ, **43**, 173-196). Thus, the total irradiance I_{tot} (in DN) is

$$I_{\text{tot}} = S_{\text{tot}} \times A(\alpha) \times 59 r^2 \sum_{\text{All pixels}} DN \left(\frac{1000\text{ s}}{t_{\text{exp}}} \right) \left(\frac{r}{1000\text{ km}} \right)^2$$

where S_{tot} is the solar conversion constant (See section V), *A* is the albedo as a function of the solar phase angle α, *r* is the range to the comet in kilometers, t_{exp} is the exposure time in milliseconds, and *DN* is the measurement of the pixels on the nucleus.

Calibration is supplied by giving the output of the Sun in dn/s at 1AU as seen through our camera and filter (see below). That number is 2.8808 x 10^{16} dn/s when calculated using 5nm steps. Further calibration to obtain absolute irradiance, including the derivation of S_{tot}, is described in Section V.

IV. System Calibration Operations

Our original intent was to take calibration images of standard star fields about every six months to determine the state of the camera and to image standards a very short time immediately before and after the encounter with Wild 2. After the contamination problem was discovered, we knew pre-launch calibrations were not likely to be of value, and it seemed particularly important to take calibration fields immediately before and after the encounter.

As optical navigation sequences were acquired as often as five times a day during final approach to encounter, it occurred to the engineers that this practice was somewhat worrisome, since electronic equipment usually fails during power cycling, so it was decided to leave the camera electronics turned on. The camera optics are bolted to the top of the electronics box, as is the CCD detector housing. As the electronics remained energized, the detector got warmer and warmer, and it soon was displaying images with high noise levels, full of hot pixels, thousands of them. The scheduled first calibration sequence provided the noisiest of all the images, and these were completely unusable. At that point it was decided to go back to power cycling, since the navigators were beginning to be bothered by the noise as well. The only useful calibration sequence was therefore the one acquired 11 days after the encounter.

The camera carries a calibration lamp. These "grain-of wheat" lamps, when operated at very low voltage, had never been known to fail and in lab tests always remained stable over many years. The calibration images of the lamp are anything but flat, of course, but we had a large number of laboratory sequences comparing the lamp to the flat field furnished by the integrating sphere. This would provide a two step calibration, current response to lamp and lamp to absolute. There was a very high-energy solar flare in November of 2003. When a lamp image was taken during a camera test a few days after the flare, the lamp apparently failed to turn on. The image was blank. At that point there was concern about trying it again, concern that there might be a short somewhere related to the apparent lamp failure, so no "flat" fields were acquired. After the encounter and the post-encounter calibration sequence, another lamp image was undertaken. The lamp worked perfectly! The only explanation that seemed to account for this was that the solar flare had flipped a bit in one of the logic circuits and the lamp was never turned on!! The best laid plans of mice and men, Etc.

The post-encounter calibration sequence was acquired at a scan mirror angle of 24 degrees without changing the attitude of the spacecraft. There was no desire to risk another off Sun and off Earth maneuver to acquire a true standard field. The ad hoc field gave us one nicely centered star of V magnitude 6.15 (SAO138420). This field was imaged seven times using three one second, three five second, and one 15 second exposure. The one second exposures were meant to prevent saturation of images of any really bright stars. (There in fact was one such star, right at the edge of the field, a star that unfortunately could not be used because it was partly over the edge in many images.) Co-adding the four longer exposures gave a reliable dn level in our instrument system for SAO138420.

V. Absolute Calibration

The throughput of the camera system was determined by measurement of curves of the quantum efficiency of the detector, the transmission of the optical system and filter, and the reflectivity of the scan mirror as a function of wavelength. For each 5 nm of wavelength, the product of these factors was tabulated to determine the transmission of the system. (See Appendix II) The reflectivity of the two periscope mirrors was also measured to be included whenever appropriate. The electronic dn (digital number) count was set at 0.05 per electron produced in the detector. The solar spectral irradiance was taken from Neckel & Labs (Solar Physics, **90**, 205-258, 1984) and convolved with the STARDUST camera parameters to determine the dn/s to be expected from the Sun at a distance of 1AU. The same parameters were determined for light passage through a Johnson V filter as well as our so-called navigation (nav) filter. This V output was converted to magnitude, and the result was –26.66, very close to the value usually quoted for the Sun (-26.74). If we had been able to use 0.5 nm wavelength intervals for these calculations instead of 5nm, the agreement would very probably have been better, but all of the data needed were not available for smaller intervals. The 0.08 magnitude difference implies our V is a factor of 0.929 of what it should be. Our intrinsic system passband is considerably wider than the V passband, mainly adding a long red tail to V. All results for both V and our intrinsic system were divided by the 0.929 factor, which should improve our absolute results.

The theoretical output of our calibration star was calculated from the parallax, spectral type, luminosity class, and V magnitude given in the Simbad database. These quantities are not sufficient to give precise values for radius and effective temperature of a particular star, of course. The parallax, radius, and effective temperature were juggled within permissible uncertainties in the parallax, spectral type and luminosity class to force an **exact** fit to the V magnitude. The values from Simbad for SAO138420 are spectral type F7 and luminosity class V, with a parallax of 29.26 ± 0.99 mas. The V magnitude is 6.15. The values used in the calculations were parallax 27.50 mas, effective temperature 6400 K, and radius 1.640 times solar. The forced fit in V guarantees that the errors in the three input quantities come close to canceling out, only close because the nav filter throughput is considerably broader than that of Johnson V to which the fit was made. The inputs given above say that SAO 138420 should produce 10230 dn/s, or 11012 dn/s when corrected, through our nav filter. When measured, the star produced 8082 dn/s (or 8700 dn/s when corrected). This suggests that our camera system then was operating at 79% efficiency overall as compared to measurements in the laboratory before launch.

We know there was still contamination present in the optical train. The reflectivity of the scan mirror could very well have decreased after exposure to vacuum, UV, etc. for six years. The quantum efficiency of the detector could have changed (we experienced a number of severe solar flares), and the calculated output of our standard star easily could be off by 10%. Given the overall environment in which the camera components existed for six years, retaining 79% of the original efficiency seems fairly good.

Lacking data from our calibration lamp, we cannot determine the system efficiency by two-step comparison with laboratory data (observed to lamp and lamp to laboratory). Given where the calibration lamp is mounted, just in front of the first element of the lens, no account could be taken of changes in the scan mirror in any case. The solar output through our optical system and nav filter, at 1 AU heliocentric distance, should be 3.101×10^{16} dn/s at full pre-launch efficiency and corrected for finite wavelength steps. The entrance pupil diameter of the STARDUST camera is 58.5 mm with a focal length a bit under 202 mm. The total measured solar input through the nav filter, including the camera throughput, pre-launch, should have been 168 W/m^2 and was 156 W/m^2 (the 0.929 factor again). Since the entrance pupil area is only .002734 m^2, These numbers imply that we actually would have an input of 0.459 W from the Sun at 1AU or 1.48×10^{-17} W/dn.

Using the information regarding the solar flux, the images are calibrated by converting the measured DN of a pixel to radiance units. The solar flux at 1 AU, integrated through the nav filter is 156 W/m^2, which corresponds to the 2.8808×10^{16} DN/s (at 1 AU) given in Section III. Each pixel has an angular resolution of 5.96×10^{-5} rad, which corresponds to a solid angle of 3.55×10^{-9} sr. So the conversion from DN to radiance is given by:

$$(156\ \mathrm{W/m^2}) / (2.8808 \times 10^{16}\ \mathrm{DN/s}) / (3.55 \times 10^{-9}\ \mathrm{sr}) = 1.525 \times 10^{-6}\ \mathrm{W/m^2/sr/(DN/s)}$$

To calibrate an image, it is first normalized by the exposure time in seconds, and then multiplied by this calibration constant to produce radiance values of $W/m^2/sr$.