Chapter Four: Flat fielding the SNAP focal plane

Michael Richmond

June 22, 2006

1. Requirements and rationale

In conventional CCD cameras, "flatfields" make corrections for the pixel-to-pixel variations in sensitivity across a detector. The SNAP focal plane, with its many detectors, is larger and more complicated than ordinary cameras, and so involves additional sources of error. We use the term “flatfielding" to mean "correcting all types of systematic photometric errors across the SNAP focal plane." We also include in this document the additional photometric effects which may occur in the spectrograph at the center of the SNAP focal plane.

In the ordinary "mowing" observational mode, the SNAP spacecraft will measure supernovae repeatedly, but always on the same set of detectors. There could very easily be small, but significant, systematic differences between, say, supernovae detected in column 3 of the East Quadrant and those detected in column 5 of the East Quadrant; or between those measured by column 3 of the East Quadrant and those measured by column 3 of the North Quadrant. These systematic differences, if not corrected, could become a dominant source of error in the cosmological calculations. It is therefore vital that we include in the SNAP operational plan special observations which can be used to characterize the several types of systematic error.

2. Imager: using internal sources

The variation in response across the SNAP focal plane, or the “flat-field”, is typically characterized by two components: small-scale pixel-to-pixel variations, or "high-frequency spatial flats", and large-scale variations, or "low-frequency spatial flats". High spatial frequency variations are usually introduced by individual pixel response differences and by shadows created by particulates deposited on filters or the CCD . To correct for these pixel-to-pixel variations, the focal plane is often illuminated with diffuse light that is spatially uniform, or "flat", so that the individual pixel responses can be normalized one to the other.

In ordinary astronomical observations, the diffuse irradiance is usually provided by a lamp source reflected off the dome (dome flat) or the twilight sky (twilight flat). For SNAP, the Ring of Fire (RoF; see Scholl 2004) functions as the diffuse source of irradiance. As shown below, the RoF places a series of lamp sources around the entrance to the cold stop. These lamps irradiate a ring of diffuse reflecting material along the opposite wall of the cold stop that scatters light onto the focal plane in a well-characterized azimuthal and radial pattern. There are several advantages to such a system. By using accurately calibrated lamp sources, the flat field irradiance can be characterized to high precision, thereby yielding very accurate high spatial frequency flat fields (monitoring the irradiance for variations with photodiodes will be discussed later).

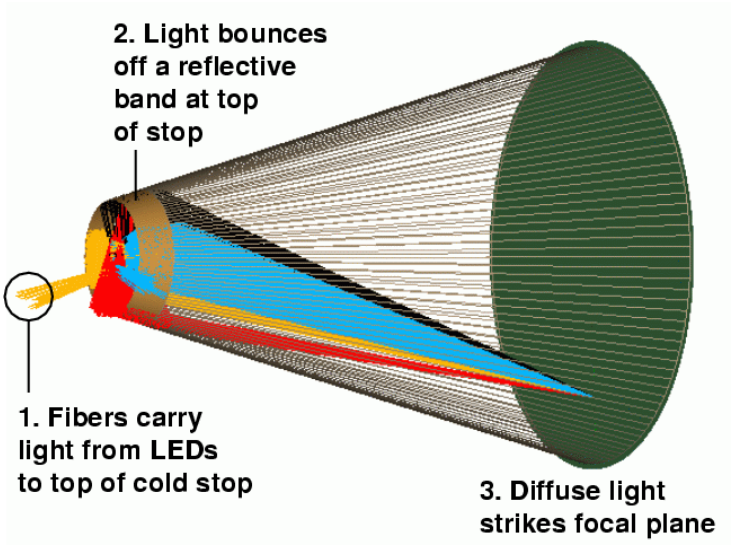

Figure 4-1 Light from LEDs bounces off the "Ring of Fire" and onto the focal plane to provide a diffuse source of light for the imaging detectors.

Similar high precision high frequency flat fields can be obtained by dithering well-characterized calibration stars over the focal plane. But the large number of pixels on the SNAP focal plane, roughly half a billion, makes this approach very costly in time. There are also other significant advantages to the RoF. First, high frequency flat fielding can be done routinely and quickly. Second, the RoF incorporates lamp sources with well-calibrated irradiance that are monitored by NIST-calibrated photodiodes. The RoF thus provides an accurate absolute flux scale for the SNAP photometry. Third, the calibrated irradiance on the focal plane from the RoF provides a means to monitor variations in the filters, as well as standard tests of detector performance.

However, the RoF does have a disadvantage when compared with standard flat-fielding techniques. While the RoF delivers calibrated diffuse light to the focal plane on small chip scales, SNAP does not currently have a scheme that puts this irradiance though the entire SNAP optical train. On the other hand, changes in the flat field due to the SNAP mirror assembly should be seen over large-scales and can therefore corrected with "star-flats" (discussed later) and with "super-flats" created from observations of the zodiacal light background.

The RoF presents some interesting challenges. Ideally, the RoF is designed to illuminate the focal plane with calibrated irradiance. Traditionally, QTH (Quartz Tungsten Halogen) lamps have been used in this role. It is not possible, however, to mount several QTH lamp sources at that position with the space constraints at the entrance to the cold stop. In addition, QTH lamps have shown that the difference in gravitational loading between ground calibration and space can change the irradiance of the lamp significantly (Martel, Hartig and Siriannai 2002; Martel and Hartig 2002). We plan to remedy this situation by using pulsed LEDs to illuminate fiber optics which deliver light to the RoF. LEDs are solid state devices that will be more stable on orbit than QTH lamps and can have controlled light output if pulsed on a low-duty cycle. LEDs also can be manufactured with a large number of narrow wavelength ranges from the optical to the near infrared which will be useful in calibrating detectors and filters.

Coupled with calibrated photodiodes on the SNAP focal plane to monitor the irradiance, the RoF will yield a steady source of irradiance that can correct small-scale flux variations, as well as monitor system response.

3. Imager: using stars

For some purposes, one can use observations of celestial sources themselves to calibrate the focal plane of a telescope, even if the properties of those sources are not known perfectly. We describe in this section how one may derive variations in sensitivity on certain scales from special measurements of stars in certain fields.

1. Intermediate-scale intra-chip variations

2. Chip-to-chip variations in QE

3. Nonuniform exposure times due to shutter

4. Filter deviations from the design

5. Changes in bandpass due to angle of incidence

6. Large-scale variations in illumination due to optics

7. Summary of observations required to characterize "flatfields"

3.1. Intermediate-scale intra-chip variations

A single detector will often have significant large-scale variations in its quantum efficiency. By "large-scale", we mean "over many tens or hundreds of pixels", or "over significant fractions of its entire extent." In addition to any which are intrinsic to the device, we will add this sort of variation when we divide images by the lamp flatfield images.

We can use "starflats" to identify these errors. The basic idea, as described by Manfroid (1995) or van der Marel (2003), is to take a series of exposures of a starfield, moving the telescope in a grid pattern so that each star is measured at many locations on a single chip. One can fit a model to the variations in observed magnitude as a function of position.

Manfroid states that a 3x3 or 4x4 grids of measurements of a field of 10-20 stars yields excellent results. On the SNAP focal plane, each detector subtends roughly 0.01 square degrees. Using counts of stars near the SNAP North field, we calculate the following cumulative statistics for number of stars falling on a single detector or filter (the optical CCDs may have 4 filters covering the quadrants of a single chip):

V mag range	***Stars per chip***	***Stars per filter (¼ chip)***
14.0 – 17.0	7	2
15.0 – 18.0	11	3
16.0 – 19.0	15	4
17.0 – 20.0	22	6
18.0 – 21.0	35	9

V mag range	*Stars per chip*	*Stars per filter (¼ chip)*
19.0 – 22.0	60	15
20.0 – 23.0	89	22

Table 4.1 Stellar density in the SNAP North field.

If we require 20 stars per chip (filter) to appear in a typical grid image, this suggests we concentrate on stars in the range from V = 17-20 (20-23). Calculations of the signal-to-noise ratio in SNAP images indicate that a star of magnitude V=20 will have S/N=100 in an exposure of roughly 100 seconds. To first order, the variations we consider here should not depend strongly on stellar color.

Required observations: a series of exposures while moving the telescope over a grid (say, 4x4 or 5x5 positions) which covers a single chip; another set of grid exposures which move stars over a single filter covering one quadrant of an optical CCD.

3.2. Chip-to-chip variations in QE

We can expect each chip to have slightly different overall quantum efficiency due to variations in the manufacturing process, especially if devices are taken from different lots. As a star moves from one detector of a given sort to another, its observed magnitude will therefore jump by some small amount.

We can determine these variations simply by moving stars from one detector to another of the same sort: that is, from an optical CCD with filter 2 to another optical CCD with filter 2. Note that this requires both relatively short offsets -- for detectors within the same quadrant of the focal plane -- and large offsets -- for detectors in different quadrants. To first order, we may treat these corrections as independent of stellar color.

Required observations: a series of exposures while moving the telescope so that stars move from one detector to all the others of the same sort.

3.3. Nonuniform exposure times due to shutter

Current designs call for four mechanical shutters near the Cassegrain focus of the telescope; see the Cassegrain Shutter document (Jelinsky 2004). Each shutter would open to allow light to reach one of the four quadrants of the focal plane. As the shutter blade rotates open, it exposes to light the inner portion of the focal plane for a slightly longer time than the outer portion. This leads to exposure times which vary across the focal plane.

Because the shutter blades move quickly (in roughly 50-80 milliseconds), this effect is significant only for short exposure times, less than 10 or 20 seconds. Jelinsky notes that it is possible to design a system to measure the motion of the blades very accurately, to within 1 millisecond, so that one could make accurate corrections with a good optical model. There are two routes one can take here:

- Calculate a correction based on the optical and mechanical design of the telescope. Call this the "theory" option.
- Determine a correction empirically, using exposures with different lengths. Call this the "empirical" option. Jelinsky (2004) and Lampton (2003, 2004) describe both methods in some detail.

It seems reasonable to do both: calculate the expected variation based on the design, and then check it once in orbit.

The effect is largest for short exposures. Consider a 1-second image: calculations indicate that stars of magnitude V=15 will yield an S/N ratio of 100-300 (highest for red stars measured on the infrared detectors). Most of the stars in the SNAP North field at this magnitude will be of spectral type G and K, which yield S/N approximately 100 in all filters; this corresponds to scatter of about 1 percent from one image to the next. Jelinsky suggests that the size of the shutter effect will be roughly 5 percent for a 1-second exposure. Thus, stars of magnitude V=15 and perhaps a bit fainter should show the effect clearly above random noise. Each shutter blade covers a single quadrant of the focal plane, which contains 18 detectors; each detector subtends roughly 0.01 square degrees, so a quadrant samples about 0.18 square degrees. In the magnitude range V=14 to V=17, we expect roughly 130 stars to be detected on each quadrant. This appears sufficient to measure the shutter effect empirically to high precision.

Required observations: a series of exposures with lengths running over a large range; say, 0.5, 1, 2, 3, 5, 10, 20, 50, 100, 200, 300 seconds. The telescope should remain fixed at one pointing during the series; it may also be possible to use a set of images with very small dithers of a few arc seconds for this purpose.

3.4. Filter deviations from the design

Although we will provide a clear specification for the SNAP filters, it is possible that small deviations may occur during the manufacturing process. Even if we measure the filters precisely before launch, it is possible that the passbands may shift somewhat after launch, or over the lifetime of the entire mission. How would these changes in effective passband affect the photometry of stars?

We may approximate such deviations from the fiducial passbands as shifts in central wavelength. Our study of passband shifts shows there is a clear pattern in the errors such shifts will produce in stellar photometry. The pattern is:

Type of star	***As central wavelength shifts***	
	Blueward	***redward***
hot, blue	grows brighter	grows fainter
cool, red	grows fainter	grows brighter

Table 4.2 Changes in apparent brightness if the central wavelength of a filter changes.

The amplitude of these changes is largest in the bluest filters of the optical CCDs and smallest in the infrared filters. Other filter changes such as a change in the bandpass width or overall transmission are more difficult to determine with stars, particularly in the infrared. As discussed in Chapter 5, we are using intermediate-band LED emission to provide additional information on filter bandpasses. Both the LED and stellar filter tracking methods can be used as independent checks of the filter transmission.

We can look for variations in passband from one instance of a given filter to another of the same filter when the telescope is first launched. We can also monitor possible changes over time

during the mission. The key here is to examine differences as a function of stellar color. Differences which do not depend on color may be attributed to chip-to-chip QE variations.

Required observations:

- a series of exposures while slewing the telescope so that stars move from one device to another with the same filter
- periodic exposures of the same star field (which will occur automatically in the "mowing" procedure

3.5. Changes in bandpass due to angle of incidence

If the SNAP filters depend on interference rather than colored glass, there will be significant shifts in the bandpass as a function of the angle with which light strikes the filters. From the inner edge to the outer edge of the focal plane annulus, this angle of incidence varies from about 0.14 to 0.28 radians. How will this affect measurements of stars?

Our analysis indicates a simple pattern that should appear as stars move radially away from the center of the focal plane: blue stars grow brighter, and red stars grow fainter. Therefore, one can characterize this effect by taking a series of images and looking at the change in instrumental magnitude as a function of stellar color and distance away from the center of the focal plane.

Required observations:

- a series of exposures in which stars move (radially) across a single filter, and (radially) from one instance of a filter to another instance of the same filter.
- Look for changes as a function of stellar color.

3.6. Large-scale variations in illumination due to optics

It is possible that the optics may cause small differences in illumination on very large scales across the focal plane; we might call such effects "vignetting." Note that such effects could not be detected using images of the internal lamps, since that light does not pass through the optics.

Another large-scale variation may come from "parasitic scattered light,'' which derives from reflections between the CCDs and the filters. While these reflections are small during normal astronomical observations, the use of bright flatfield lamps can contribute to the scattered light at all positions on the focal plane. Normalizing the science images with a flatfield contaminated by such variations will lead to position-dependent systematic errors.

Preliminary estimates are that any such effects would depend only very weakly on the wavelength of light, and hence only very weakly on stellar color. We believe that these very large-scale variations would be removed by the application of corrections already mentioned; specifically,

- Intermediate-scale intra-chip variations
- Chip-to-chip variations in QE
- Required observations: None.

3.7. Summary of stellar observations required to characterize "flatfields"

We end up with the following list of stellar exposures to determine "flatfield" corrections:

- a series of exposures while moving the telescope over a grid (say, 4x4 or 5x5 positions) which covers a single chip;
- Another set of (4x4 or 5x5) grid exposures which move stars over a single filter covering one quadrant of an optical CCD.
- a series of exposures while moving the telescope so that stars move from one detector to all the others of the same sort (both within the same quadrant on focal plane, and across to the other three quadrants)
- a series of images with exposure lengths running over a large range

We plan to make these special observations at the start of the mission and periodically thereafter.

4. Dithering and Mowing

The previous section describes a series of special calibration exposures which should be made at the start of the mission and periodically thereafter. For the bulk of the time, however, the SNAP telescope will execute its normal program of repeatedly scanning a small region of the sky. Let us consider briefly the details of this procedure: is there a "best" way to move the telescope across the sky?

There are two main issues we can address in designing the slewing procedures:

- mapping the variations in sensitivity within a single pixel of a detector, also known as intra-pixel variations; we will use term dithering to describe the telescope motions used to characterize it
- covering a contiguous region of sky completely over the course of several passes; we will use the term mowing to describe these large-scale motions

First, let us discuss dithering.

The following discussion assumes that we can control the telescope well enough to point it to a particular sub-pixel location reliably. If that is not the case, then we should simply command the telescope to move by a pixel or so between each exposure and take whatever random dithering we get. We will have enough stars on each detector to extract the exact offset of each sub-exposure from the others after the fact.

The regular observations will involve multiple exposures at each position. As HST and other space- based telescopes have shown, cosmic rays strike a significant fraction of the pixels in a detector over periods of just a few hundred seconds. In order to reach supernovae of magnitude 25 or so with decent signal-to-noise ratios, SNAP must collect light for about 1000 to 1500 seconds. It therefore makes sense to follow the practice of HST and break up the total exposure into several pieces: say, four images, each of 300 seconds. Note that at exposure times of several hundred seconds, noise from the background sky (mostly zodiacal light) will roughly equal the readout noise for optical CCDs. Readout noise will exceed the background sky noise for the current versions of infrared detectors until the exposure time reaches several thousand seconds. Finally, breaking each exposure into several pieces increases the dynamic range of the dataset, since brighter stars will be recorded without saturation on the shorter individual exposures.

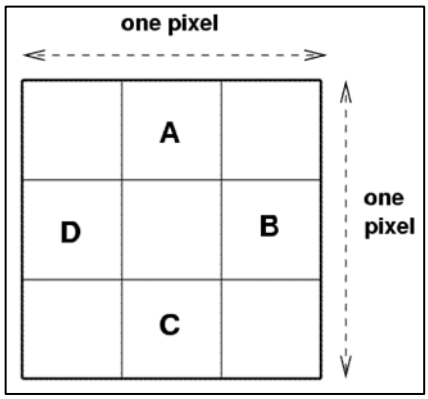

How should we dither the telescope from one sub-exposure to the next? Several scientists have studied the issues involved in sub-pixel dithering. Lauer (1999) finds that moving in a regular NxN grid-like pattern is best for photometry. Bernstein (2002) concludes that a 3x3 grid is nearly always sufficient to recover the original properties of point sources. Suppose that we make one

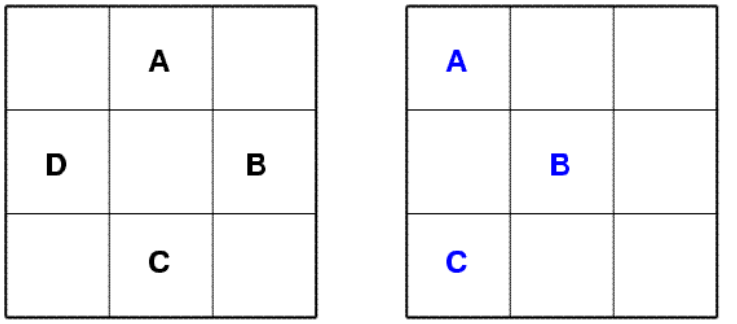

a star in pixel (100,100) another star in pixel (323,1408)

set of four exposures using small offsets of size one-third of a pixel each time:

Each individual star will sample only about half of the 3x3 intra-pixel grid during this procedure. However, there will be tens or hundreds of stars on each chip with good signal-to-noise. These stars will be scattered at random across the sub-pixel locations, so that some stars will fall into those other sub- pixel locations:

If we assume that the manufacturing process causes the same pattern of intra-pixel sensitivity across each detector, then we can use the many stars measured on each chip in a single set of four sub- exposures to sample all locations in a 3x3 sub-pixel grid.

We now consider a second aspect of telescope motion: the large-scale slews which point the telescope to all portions of the study area, which we denote as the "mowing" pattern.

Because detectors do not cover the entire focal plane, there will be small gaps in the sky coverage between each set of images. For simplicity's sake, let us illustrate the issue with just two stars, and focus on their motion relative to the top left quadrant of the focal plane during a single pass of the mowing:

If we move the telescope in a straight line across the sky, then it will miss all stars which fall in the gap between two rows of the camera. The gap is roughly 18% of the width of each detector, so each of these linear passes will detect only 82% of the stars in a region.

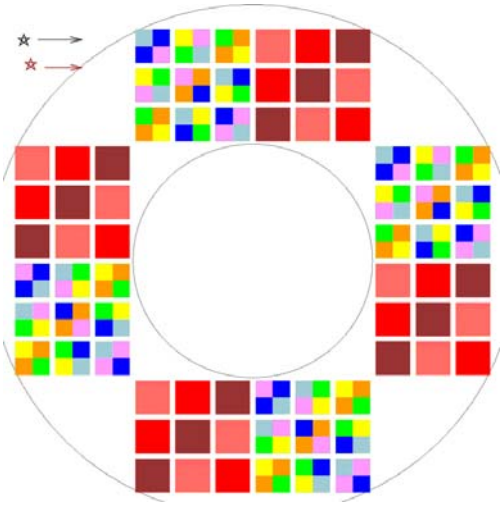

4-1 During ordinary "mowing" operations, stars move across the focal plane from left to right. The text follows two stars which will slide across the top left quadrant of detectors.

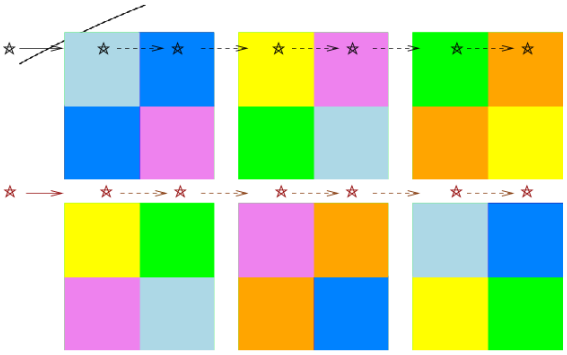

In order to cover a contiguous region, the telescope must make at least two passes. The second pass must involve a shift perpendicular to the scan direction by an amount sufficient to move stars from the inter-row gap to the detectors, like so:

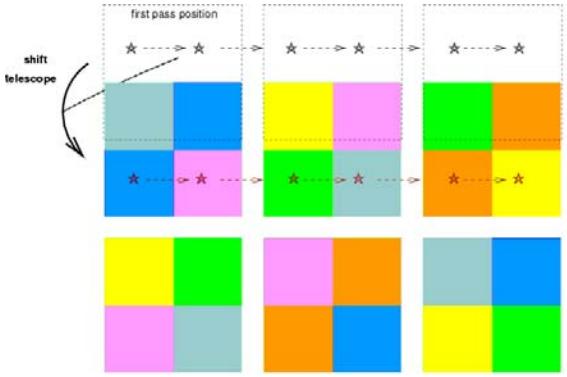

Figure 4.8 A sequence of alternating short and long moves along the scan direction causes objects to appear at the same pixel location in all detectors, but some objects will repeatedly fall between chips.

However, there is another gap to consider: the space between two detectors in the scan direction. There are several options to handle this gap. We may choose to slew the telescope along the scan direction in alternating short and long jumps, so that a given star will always appear in the same location within each filtered section of a CCD detector. This will ensure that during one pass of the telescope, many objects will be measured through every filter.

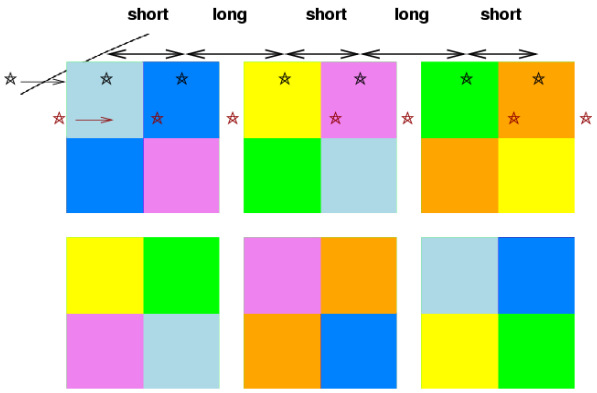

But, as the figure above shows, other objects will repeatedly fall into the gaps between detectors in a row. In order to measure every object through every filter, we must again plan a second pass, this time shifting the telescope's position parallel to the scan direction.

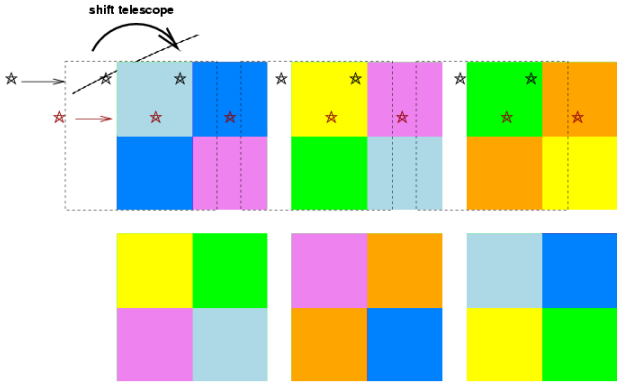

The second pass will involve an offset in both directions from the starting point of the first pass. Since the gaps between columns of detectors are roughly the same size -- about 18% of the detector size -- the yield of two passes with this staggered offset will be

- one measurement of every object in all passbands
- a second measurement of about (100-18)*(100-18) = 67 percent of all objects in the region

One way to proceed, then, is to make two passes through the survey region. Each pass will consist of pairs of alternating short and long jumps in the scan direction. After completing one pass, the telescope should return to the starting point, offset itself by a fraction of the detector size in both directions (parallel to and perpendicular to the scan), then make another series of short and long jumps as it moves down the region a second time. We suggest this pattern -- go all the way to the end of the survey region in a straight line,

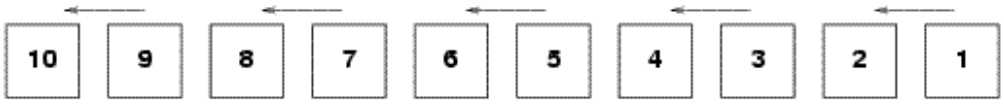

then return for a second long pass:

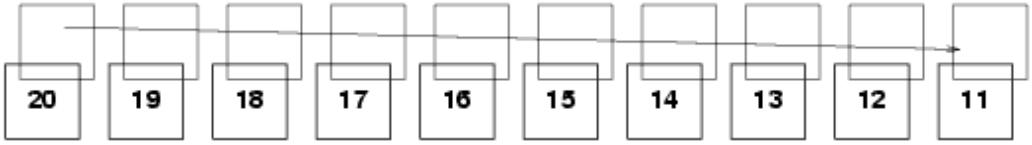

rather than a zig-zag pattern

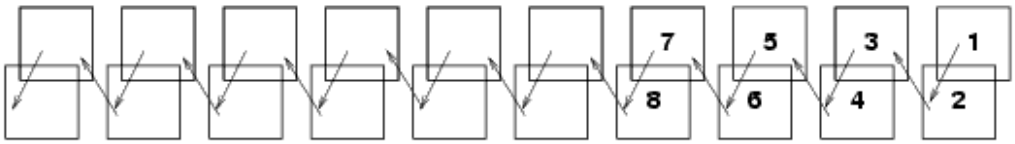

because the first method provides better time coverage of the doubly-measured objects than the second. For example, if a star falls onto the detectors in each pass, we will acquire for it

- straight path: one measurement, another 4 days later, another 4 days later, another 4 days later, etc.
- zig-zag path: one pair of measurements 20 minutes apart, then another pair 8 days later, then another pair 8 days later, etc.

Note that if it is possible to position the telescope at desired sub-pixel locations with some accuracy, we should make the offset between the first and second passes not an integer number N of pixels in each direction, but an integer plus a fraction N + 1/3. In this way, during the second pass, each sub-exposure can fall on a different sub-pixel then the sub-exposures in the first pass. After two passes, we would not only have at least one measurement of each object in the survey area through every filter, we would also place some stars on 8 out of the 9 sub-pixel locations in a 3-by-3 sub-pixel grid.

5. Algorithms

Several astronomers have considered the ways in which one can use multiple observations of stars at different locations across the field to derive the variations in sensitivity across a detector.

Manfroid (1995) treats the problem in a classical fashion. He makes a model for variations in sensitivity as a low-order polynomial of position on the focal plane, calculates the sum of squares of differences between model and measurements, differentiates the sum with respect to several parameters, and uses the zeroes of the derivatives to solve for the parameters. In other words, he employs a standard linear least-squares approach.

Van der Marel (2003) follows the same general approach, but investigates ways in which the computations can be simplified. He decreases the number of unknowns by setting aside (initially) the magnitudes of the stars in the ensemble, concentrating only on the parameters describing the change in sensitivity. This greatly decreases the size of the matrix equations and speeds up the calculations. He also shows how breaking a large detector into smaller subunits may permit one to fit realistic variations in sensitivity more accurately than with standard low-order polynomials applied across the entire field.

In our early analysis, we have so far followed the classical approach: treat both the magnitudes of field stars and properties of the detector as unknowns, and solve for all simultaneously. SNAP will look at relatively sparse fields, far from the galactic plane; during the short exposures we plan for calibration purposes, it will detect relatively few stars bright enough to have negligible photon noise. With only a few hundred to a few thousand stars serving as the sources in our calculations, we have no need to optimize our algorithms for speed or memory usage.

As an example, we provide an explicit description of one photometric equation we have used in our simulations. Consider the conversion of an instrumental magnitude m to its equivalent M on some standard system. With a perfect single detector, one would simply make a single shift:

$$M = m + a$$

where **a** is a zero-point offset term. However, there are 72 different chips in the SNAP focal plane, which makes the equation

$$M = m + a_i$$

where $\mathbf{a_i}$ is the zero point for chip *i*.

If the instrumental bandpass doesn't match the standard bandpass exactly, then there will be small corrections which depend on the color of the star.

$$M = m + a_i + b_j(color)$$

where $\mathbf{b_j}$ is the first-order color term for a particular chip-plus-filter, and *color* is some measure of the star's color.

But the effective bandpass will shift slightly across each chip-plus-detector, because the angle of incidence will change. We may need to take this into account, in which case we would need to replace the single color term with a more complicated expression:

$$M = m + a_i + b_{1\,j}(color) + b_{2\,j}\theta(color)$$

where we now have a constant $\mathbf{b_1}$ and a slope $\mathbf{b_2}$ term for each chip-plus-detector, and θ is the angle of incidence at which light strikes the detector.

If the sensitivity of each chip is not perfectly uniform across its face, then we need to correct for this 'small-scale" flatfielding error. We might approximate the changes in sensitivity as a low-order polynomial function **p** of (*row*, *col*) position on the chip.

$$M = m + a_i + b_{1\,j}(color) + b_{2\,j}\theta(color) + p_{1i}(row) + p_{2i}(row)^2 + p_{3i}(col) + p_{4i}(col)^2$$

Note a significant difference between these simulations and our experiments with real data. The real data consists of measurements made with a single CCD which is centered on the optical axis of its telescope; as a result, the large-scale variations are radially symmetric around the center of the chip. The simulated data, on the other hand, come from detectors scattered all over the SNAP focal plane, all of which are far from the optical axis. We therefore expect asymmetric patterns in sensitivity across each detector.

We have not yet put the color-dependent terms ($\mathbf{b_1}$ and $\mathbf{b_2}$ in the equations above) into our simulations and analysis.

It is likely that after SNAP has gathered months or years worth of measurements of stars in its ordinary operations, we will want to analyze this much larger collection of relatively faint stars to look for small, uncorrected systematic errors. In that case, it might well be prudent to follow the methods set forth by van der Marel (2003) in order to speed up the computations.

6. Testing our procedures: observations and simulations

We are testing this general approach to characterizing variations in sensitivity in two ways: by

making repeated observations of star fields with a real telescope, and by running artificial data through a simulation of the SNAP telescope. Let us describe briefly our analysis in each case.

6.1. Observations with the WIYN telescope

We have undertaken an observing program at the WIYN 0.9m telescope to test the stellar flat-field method using the single-chip S2KB imager. We were mainly interested in the remaining error level after the stellar flat correction is applied as well as developing an efficient, repeatable observing technique. Following suggestions from Manfroid (1995), we chose targets among the Stetson Cluster Standards that would sufficiently populate our field of view without too much overlap of stellar sources. For the S2KB, the sparse cluster of NGC2420 fills about 1/3 of the CCD and tracks well over zenith to reduce differential reddening within the field. Our observing cadence pointed the cluster to a 3 by 3 grid covering the entire CCD in a single filter band. In some cases, we took multiple exposures per readout of the CCD to increase our efficiency of observing time.

To pull out the residual spatial differences in photometry across the chip, we first calculate the mean instrumental magnitude of each star with less than 1% Poisson noise. We then calculate the difference of each instrumental magnitude measure of a single star from its mean magnitude as a function of position on the CCD. This difference can be weighted by the poisson noise in a χ^2 fit by a spatial correction function to the residuals. The formula is:

$$\chi^2 = \frac{(mag(x,y,s) - mean\_mag(s) + CF(x,y))^2}{\sigma(s)^2}$$

where **s** is a single star, **x** and **y** are the star's coordinates on the chip, and ***CF*** is the spatial correction function. As employed by Manfroid and Van der Marel, the correction function is typically a low order 2D polynomial to match any spatial residual while remaining well behaved at the boundaries of the chip. In our tests of S2KB, we found a strong central residual on the chip of about +0.04 magnitudes which could not be well fit by a low order polynomial.

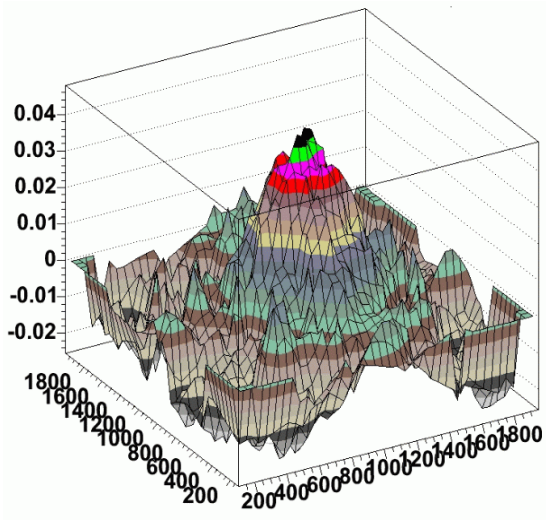

Figure 4.13 Residuals in photometry (magnitudes) as a function of position on the detector (pixels) for data taken at WIYN.

Therefore, we used 2-D "Penny2" function (gaussian core with lorentzian wings) which could well fit the central peak, maintain continuity at the chip edges, and keep a low number of fit parameters. We minimize χ^2 with this function and determine the fit parameters.

Once the spatial correction function has been calculated, we first applied it to the all of the stars that had less than 0.3% Poisson error in their photometry. Since these same stars contributed most heavily to the χ^2, this essentially gives us a measure of how well we fit the data. Our results show that the final residual in the fit of our stellar flats is about 0.005 mag RMS. The true test of how well this residual corrects the flatness of photometry across the chip is by applying the

correction to another set of stellar measurements dithered over the entire chip. Again using stars with 0.3% statistical photometry error, our final spatially-corrected residual error degrades slightly to 0.006 mag.

7. Simulations of the SNAP focal plane

We have written a self-contained package for simulating various effects involving variations across the SNAP focal plane. The code is written in a mixture of TCL and C, and is freely available at

http://spiff.rit.edu/richmond/snap/pipeline/aug19_2004/snap_pipeline.html

This is not a pixel-level simulator; its basic units are stellar measurements. The user provides an input set containing

- a catalog of stars with known spectral types and magnitudes
- a description of the focal plane (how many detectors, their locations and response functions, etc.)
- a prescription for the type of systematic errors to be included in the run (i.e. zeropoint offsets for each chip, polynomial function of sensitivity variations across each chip, etc.)
- a recipe for observing (exposure times, dithers)

The program carries out a set of observations, calculating the magnitudes which would be produced in each image from each chip. Although the calculations do not treat individual pixels in each detector, they do include the effects of photon noise, dark current, and other effects on a statistical level.

To illustrate the purpose of this photometric simulator, we show below an example of one set of tests which were made with it.

The input stars in this simulation are based on the USNO-A2.0 catalog of the Northern SNAP field. We point the telescope at RA=270 degrees, Dec=+67 degrees, and then consider only the eastern quadrant of the detector. We include all stars with B and R magnitudes brighter than 18.0, in a one-degree box around the center of the eastern quadrant (RA=271.28, Dec=67.0). That yields about 1840 stars. We assign a spectral type to each star based on its (B-R) color from the USNO-A2.0 catalog.

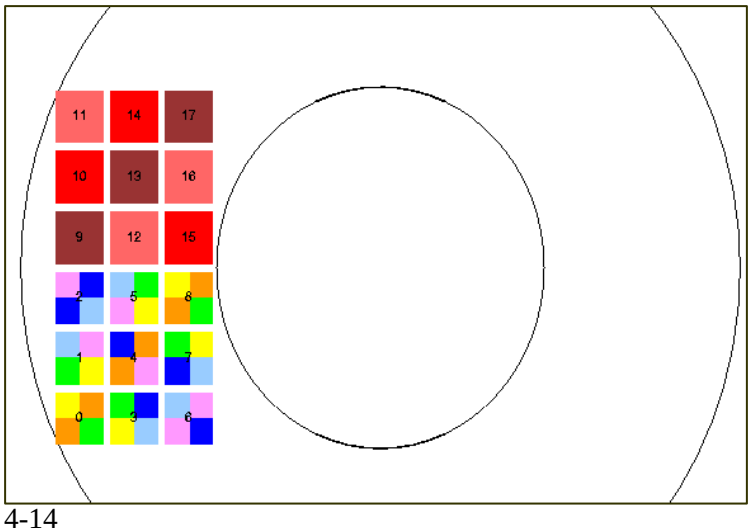

Figure 4-2 A realistic model of the East Quadrant of the focal plane. The CCD chips are each covered by four optical filters, while each IR chip has a single filter.

We used both a simplified, "monochromatic" version of the focal plane, in which all detectors were optical CCDs with fiducial filter 5, and a realistic focal plane, with a block of optical CCDs and a block of near-IR detectors:

We moved the simulated telescope in a 6x6 grid-like pattern, so that some stars would move across one entire block, appearing at least once on each visible CCD, or at least once on each near-IR detector.

All exposures were 10 seconds long. We used two modes of observing:

- take one image (snapshot) per position; total 6x6x1 = 36 images
- take 4 images per position; total 6x6x4 = 144 images

Of course, in any particular snapshot, some stars will fall between detectors. Our analysis used Honeycutt's inhomogeneous ensemble photometry technique, including any stars which are detected on at least 10 images.

We did NOT include any of these complicating effects:

- intra-chip sensitivity variations
- shift in bandpass due to angle of incidence
- color terms between supposedly identical filters
- shutter effects

The simulator can add all these effects to the "observations", but we are not yet ready to analyze the results in an automated fashion.

In this preliminary simulation, we checked to see how well one can determine the chip-to-chip offsets from a set of exposures which move across a field of stars. We find

- with a monochromatic focal plane, a 6x6 grid determines chip-to-chip offsets to less than 1 mmag
- with a realistic focal plane, a 6x6 grid of 1 snapshot at each position determines offsets to a precision of 1-7 mmag (depending on the filter)
- and a realistic focal plane, a 6x6 grid of 4 snapshots at each position determines offsets to a precision of 1-2 mmag (depending on the filter).

8. References

Bernstein, PASP, 114, 98 (2002) Advanced Exposure-Time Calculations: Undersampling, Dithering, Cosmic Rays, Astrometry, and Ellipticities

Bonissent, Ealet, Smadja (SNAP Technical Memo, 2004) Spectrograph calibration short exposure times

Ferruit (SNAP Technical Memo, 2003) Calibration of the SNAP IFU spectrometer: error budget

Honeycutt, PASP, 104, 435 (1992) CCD ensemble photometry on an inhomogeneous set of exposures

Jelinsky (SNAP Technical Memo, 2004) Cassegrain shutter concept/proposal

Lauer, PASP, 111, 1434 (1999) The Photometry of Undersampled Point-Spread Functions

Manfroid, A&AS, 113, 587 (1995) Stellar calibration of CCD flat fielding

Martel, Hartig and Sirianni (ACS Technical Memo, 2002) On-orbit HRC and WFC Internal Tungsten Lamp Count Rates

Martel and Hartig (ACS Technical Memo, 2002) HRC and WFC Internal Tungsten and Deuterium Lamp Count Rates

Scholl (SNAP Technical Memo, 2004) Options for Placement of Flatfield Sources

van der Marel, STScI ISR ACS-2003-10 (2003) Determination of Low-Frequency Flat-Field Structure from Photometry of Stellar Field