

Characterizing SNAP NIR Photodetectors with the Spot-O-Matic

Nathaniel Barron

Advisor: Wolfgang Lorenzon

A Thesis presented for the degree of
Bachelor of Science

SNAP Group
Department of Physics
University of Michigan
April 2006

Characterizing SNAP NIR Photodetectors with the Spot-O-Matic

Nathaniel Barron

Submitted for the degree of Bachelor of Science
April 2006

Abstract

The SNAP project (SuperNova Acceleration Probe) will be a space-based telescope, which will capture images with a focal plane array of 36 CCDs as well as 36 HgCdTe detectors for near infrared light. The University of Michigan contingent of the SNAP collaboration is responsible for testing the HgCdTe detectors and ensuring that they meet the SNAP performance requirements.

The University of Michigan team receives near infrared devices from two manufacturers competing for a contract to deliver 36 science-grade devices. Devices from both manufacturers are characterized in the lab, suggestions for fabrication improvements are returned, and eventually a manufacturerer will be chosen based on its ability to meet SNAP science requirements.

The following characteristics of the devices must be fully understood and documented:

- Gain
- Dark Current
- Noise
- Response linearity
- Quantum efficiency
- Pixel response function
- Inter-pixel uniformity
- Lateral charge diffusion

This thesis documents the creation of an experimental apparatus designed to fully characterize the pixel response function, inter-pixel uniformity, and lateral charge diffusion of the SNAP detectors. It will also present experimental results obtained from an InGaAs test device made by one of the competing manufacturers, Raytheon Vision Systems. This test device shares many characteristics with the actual HgCdTe devices that are to be used. In addition to providing valuable information about the test device itself and the manufacturing process, it serves as a proof of functionality for the experimental apparatus.

Declaration

The work in this thesis is based on research carried out by the Michigan SNAP group, in the Department of Physics at the University of Michigan, Ann Arbor. No part of this thesis has been submitted elsewhere for any other degree or qualification and is all my own work unless referenced to the contrary in the text.

Copyright © 2006 by Nathaniel Barron.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

Acknowledgements

I would like to thank the following people for their invaluable contributions to my research and to the completion of this thesis:

- Wolfgang Lorenzon, my advisor, for his guidance, insight, and limitless patience. I can't say enough about the endless support he has given me.
- Michael Borysow, who wrote all of the motion control and analysis software with which the measurements contained in this thesis were carried out. Without him, none of these results would have been possible.
- Andrew Tomasch, for teaching me the tricks of the trade, and just about everything I know about hardware.
- Greg Tarle, for trusting me with his time and money, and for always seeing the light at the end of this very long tunnel.
- The entire Michigan SNAP group, including Matt Brown and Michael Schubnell, for their teaching and intellectual guidance.

Contents

Abstract	ii
Declaration	iii
Acknowledgements	iv
1 Introduction	1
1.1 The SNAP Telescope	1
1.2 SNAP Detector Requirements	1
1.3 Experimental Principle	3
2 Experimental Setup	5
2.1 Spot-O-Matic Optics	5
2.2 Experimental Setup	10
3 Experimental Procedure	11
3.1 Characterizing the Spot Size	11
3.2 Focusing the Spot on the Device	13
3.3 Tests to be Performed	15
4 Characterizing the Pixel Response Function	18
4.1 Discussion of the Pixel Response Function	18
4.2 Extraction of the Pixel Response Function	19
5 Photometry Requirements	24
5.1 Initial Search for Dips	24
5.2 Non-linear Response	25
5.3 Multiple Pixel Scans	26
5.4 SNAP Airy Disk Spot Scans	29
5.5 Dead Pixels	32
6 Improvements	35
6.1 Encoder Backlash	35
6.2 Light Source Stability	36

7 Further Tests with the Spot-O-Matic	40
7.1 Rockwell 1k banded HgCdTe Array	40
7.2 Rockwell H2RG HgCdTe #40	43
7.3 Rockwell H2RG HgCdTe #102	45
Bibliography	48
Appendix	49
A Encoder Mounts	49

List of Figures

1.1	Quantum efficiency as a function of wavelength for several different operating temperatures for the Raytheon Vision Systems (RVS) SB-301 InGaAs detector.	4
2.1	Measured magnification and working distance as functions of the distance between the pinhole and the tube lens.	6
2.2	Diagram of the optical components of the Spot-O-Matic.	7
2.3	Sample light intensity over time of Photomax light source, with and without feedback regulation. Figure courtesy of Oriel Instruments. [5]	8
2.4	Histograms of backlash results for the x and y axes of our stage. It can be seen that not only is the backlash as large as several microns at times, but it is also not constant, making it difficult to compensate for it in software. Graphs courtesy of Michael Borysov [7].	9
2.5	The fully functional Spot-O-Matic, positioned in front of the dewar window, ready to perform measurements on the InGaAs detector.	10
3.1	Diagram of knife-edge scanning procedure.	11
3.2	Knife-edge scans for $1050nm$ and $1550nm$ light at focus. The top panels show intensity recorded on the detector as a function of spot position relative to the edge of the razor blade. The bottom two panels show the derivative of this intensity plot, which gives us the intensity profile of the spot.	12
3.3	A plot of σ versus focus position for $1050nm$ light. The black lines represent the variation of spot size with distance far from focus. Near the focus, the spot size changes more like a quadratic function, shown with the red curve. The minimum σ is $1.40\mu m$. The profile of this spot is plotted in Fig. 3.2.	13
3.4	A virtual knife-edge scan through a sample pixel at optimal focus with $1550nm$ light.	14
3.5	A plot of σ versus z -axis position from virtual knife-edge scans performed on a sample pixel at $1550nm$.	15
3.6	Schematic view of a pixel scan. The red line shows the motion of the spot, and the dots represent places where the stage stops and an image is taken.	16

4.1	The top two panels are raw FITS files of a small section of pixels illuminated at each wavelength, viewed with the SAOImage DS9 software. The difference in the response of the neighboring pixels is slight, but noticeable. The bottom panels are histograms of the pixel read-out, normalized, and with the range adjusted to highlight the difference in lateral charge diffusion. The diffusion is slightly higher for $1050nm$	21
4.2	Single pixel scans at $1050nm$ and $1550nm$. The grid on the bottom represents the physical size of the pixel ($20\mu m$).	22
4.3	One-dimensional virtual knife-edge scans through the center of a pixel, at $1050nm$ and $1550nm$. The derivatives of the data points are shown below, with the Gaussian fit of the peaks in blue.	23
5.1	Virtual knife-edge profiles for 3 adjacent pixels at both wavelengths. The response of all 3 is summed at each point and plotted above the individual profiles.	25
5.2	The top panel shows the sum of 9 adjacent PRFs scanned at $1050nm$ with the small diffraction-limited spot in Fig. 3.2. The response histogram is shown below.	27
5.3	The top panel shows the sum of 9 adjacent PRFs scanned at $1550nm$ with the small diffraction-limited spot in Fig. 3.2. The response histogram is shown below.	28
5.4	The top two panels show the knife-edge scan spot profiles created with the $1000\mu m$ pinhole at $1050nm$ and $1550nm$. The red curve is the calculated Airy disk function for the SNAP telescope at the respective wavelengths. The bottom panel compares the shape of the measured spot profiles at each wavelength. This demonstrates that the discrepancy in fits in the top panels is due to the wavelength dependence of the calculated Airy disk, as described in Section 5.4.	30
5.5	The top panel shows the sum of 9 adjacent PRFs scanned at $1550nm$ with the SNAP-like large spot in Fig. 5.4. The response histogram is shown below.	31
5.6	FITS image of the dead pixel to be scanned.	32
5.7	Response vs. spot position for an assumed dead pixel at $1550nm$	33
5.8	Summed pixel scan of the area around an assumed dead pixel at $1550nm$	34
5.9	FITS image of a region on the InGaAs with several “daisies,” illuminated with background light.	34
6.1	A single pixel scan with a measurement error caused by stage backlash.	36
6.2	Full 5×5 summed pixel scan at $1050nm$. Fig. 5.2 shows a subregion of this scan. The response histogram is shown below.	38
6.3	Full 5×5 summed pixel scan at $1550nm$. Fig. 5.3 shows a subregion of this scan. The response histogram is shown below.	39
7.1	Single pixel scan from the Rockwell 1k banded HgCdTe Array at $1550nm$	41
7.2	Summed pixel scan from the Rockwell 1k banded HgCdTe Array at $1550nm$	42
7.3	Summed pixel scan from the Rockwell H2RG HgCdTe #40 device.	43
7.4	Single pixel scan from the Rockwell H2RG HgCdTe #40 device.	44

7.5	One-dimensional slice of a single pixel scan from the #40 device.	45
7.6	The top panel shows the summed pixel scan from the #102 device at 1050nm. The response histogram is shown below.	46
7.7	The top panel shows the summed pixel scan from the #102 device at 1550nm. The response histogram is shown below.	47
A.1	Diffraction Grating Scale Mount	50
A.2	Encoder Head Mount	51

Chapter 1

Introduction

1.1 The SNAP Telescope

The SNAP mission seeks to map our night sky with unprecedented detail, and in doing so, help us to better understand the acceleration of our universe and provide insight into the nature of the mysterious dark energy. SNAP will find and analyze thousands of type Ia supernovae and use them as standard candles to map the expansion history of the universe. SNAP will be able to map much more of the universe than existing telescopes, with a field of view of 1 square degree, 600 times that of the Hubble Space Telescope. It will also be able to see supernovae in the crucial NIR spectrum, with redshift coverage of $z = 0.4$ to $z = 1.7$.

SNAP will be a three mirror, 2 meter, f/10 reflecting telescope. The focal plane will consist of 36 $2k \times 2k$ $10.5\mu m$ pixel CCDs and 36 $2k \times 2k$ $18.0\mu m$ pixel HgCdTe detectors. The CCDs are sensitive from $\lambda = 0.4\mu m$ to $1.0\mu m$, and the HgCdTe from $\lambda = 1.0\mu m$ to $1.7\mu m$.

1.2 SNAP Detector Requirements

The SNAP mission places stringent requirements on the characteristics of these detectors. In order to achieve the precision photometry required by SNAP, the response of the devices must be uniform over the entire device to within 2%. That is, the quantum efficiency of the device must be the same everywhere, to within 2%. There are three basic problems which could hinder this requirement. First, the response within a single pixel might not be uniform. This is known as intra-pixel variation. Second, it is possible that quantum efficiency could be lower between pixels. This expected dip would be very small and hard to detect. Third, there could be overall sensitivity differences among different individual pixels. This is commonly fixed with a flat-field correction and so is not as crucial as the other two. These latter two device features are known as inter-pixel variation.

The gaps between pixels and non-uniformity within a single pixel are really two manifestations of the same problem. This problem arises because of the large size of the pixels on the HgCdTe detectors. These pixels will be $18\mu m$ square, almost twice the size of the CCD pixels on the telescope. Since the junction is attached to the substrate below the center of the pixel, it is not unreasonable to suspect that electron-hole pairs created farther away from

this junction might be collected with less efficiency than those pairs that are created close to the junction. Therefore, we might expect regions of lowered sensitivity near the edges and corners of pixels, and at the pixel boundaries. We are concerned about any such areas that would cause information to be lost or obscured. We expect these problems to also manifest themselves in the Raytheon InGaAs test device, since it shares many of the aforementioned characteristics with the HgCdTe devices.

In addition to the SNAP photometry requirements, measurement of the pixel response is crucial when we consider the problem of undersampling. Since the HgCdTe pixels are so large, the raw images obtained by SNAP will be undersampled. The Airy disk produced at the focal plane by the telescope optics has an angular intensity profile given by

$$I(\theta) \propto \left[\frac{J_1(\beta)}{\beta} \right]^2, \quad (1.1)$$

with β given by

$$\beta = \frac{a\pi}{\lambda} \cdot \sin \theta, \quad (1.2)$$

where a is the aperture diameter. The SNAP telescope will have a $2m$ aperture and $20m$ focal length. Thus the smallest Airy disk that can be produced by the telescope optics will have a FWHM of 0.105 arcsec at $\lambda = 1000nm$ (the low-end cutoff for HgCdTe devices). This corresponds to a width of $10.2\mu m$ at the focal plane, while the HgCdTe pixels are $18\mu m$ across ($20\mu m$ for the InGaAs test device). In order to be Nyquist sampled, the pixels should be no larger than half as wide as the FWHM of the Airy disk produced at the lowest wavelength. [1] By this standard, SNAP is severely undersampled.

In order to recover some of this lost information, SNAP will employ a 2×2 dither pattern when capturing images. This will greatly improve the resolution of the telescope, and allow us to recover much of this lost information, though it still will not quite be Nyquist sampled. To be truly Nyquist sampled, SNAP would need to employ a 4×4 dither pattern. However, for photometry measurements of time-variable point sources (supernovae, for example), a 2×2 dither pattern is an adequate compromise between speed of measurement and information loss. [2]

A 2×2 dither scheme has the ability to effectively cut the size of the pixels in half, but only if the properties of the telescope optics and the detectors are known to high accuracy. Standard strategies for reconstructing dithered images, such as Drizzle [3], require accurate knowledge of the point-spread function (PSF) of the optics, the size of the physical pixel, and the response function of the electronic pixel, or pixel response function (PRF). The PRF is a two-dimensional map of a single pixel's electronic response as a function of a photon's striking position inside or outside the physical area of the pixel. For Drizzle in particular, the observed image is represented as a convolution of the "true" image with four other functions. This convolution is described by

$$I = T \otimes O \otimes E \otimes P \otimes G, \quad (1.3)$$

where I is the observed image, T is the "true" distribution of light in the sky, O is the PSF of the optics of the telescope, E is the PRF of the pixels, P is the physical size of the pixel, G is the size of the finer output grid of "virtual" pixels, and \otimes is the convolution operator.

In order to extract the “true” distribution of light on the sky, T , we must accurately know all the other five functions. I is what we observe, O is known very accurately for the SNAP telescope, P is known very accurately from the manufacturer, and G we determine based on our reconstruction scheme. Therefore, the only unknown function is E , the pixel response function, which has never been measured for any of the SNAP detectors. Once it is measured, we will have all the tools to successfully dither and reconstruct astronomical images, thereby overcoming our problem of undersampling. Once the PRF is measured, it will also tell us about any loss of charge inside of the pixel, and by summing up neighboring PRFs, we will see any loss of charge between pixels. Thus finding the PRF becomes our main goal, since it will allow us to verify the SNAP photometry requirements as well as reconstruct dithered images.

A brief discussion of the PRF is in order. A more thorough discussion awaits in Chapter 4. It may help to flip quickly to page 21 to get a mental picture of what the PRF looks like. For an idealized pixel, the PRF would be a two-dimensional square function with width equal to the physical pixel width. This means constant electronic response for a photon striking anywhere within the pixel, and zero response for any photon striking outside the physical area of the pixel. As page 21 clearly shows, real pixels do not look like this. They will deviate from the square-function ideal for a number of reasons, including pixel defects. Most notably, the PRF is non-zero outside the physical pixel area, for reasons of lateral charge diffusion that will be described in Chapter 4. Thus it is crucial that we accurately map the PRF for any device we are testing.

It turns out that the PRF of the device also depends on the wavelength of incident light. Therefore, all measurements presented in this thesis are conducted at two different wavelengths, $1050nm$ and $1550nm$. These wavelengths are chosen because they reside at opposite ends of the region of highest QE for the InGaAs test device, and are separated enough to allow wavelength-dependent effects to become visible. The quantum efficiency vs. wavelength for the InGaAs detector is shown in Fig. 1.1, for several different operating temperatures. The HgCdTe detectors are responsive over a slightly larger range, $1000nm$ to $1700nm$, so any future tests on these devices will have to be conducted over this wider range.

1.3 Experimental Principle

While measuring any large-scale quantum efficiency variations across the device is relatively easy (a simple flat-field would do), performing the small scale measurements such as mapping the PRF and searching for the expected dips between pixels is not as easy. When working with either $20\mu m$ InGaAs pixels or $18\mu m$ HgCdTe pixels, any dips between pixels or regions of reduced sensitivity in the corners would be no larger than a few microns. In order to resolve structure in the PRF on the micron level, the pixels must be probed with a spot of light of similar size.

A common solution to the problem of probing the response of semiconductors is the pinhole projector. The basic idea of a pinhole projector is relatively simple. Light is shone through a tiny circular pinhole, which is placed at the focal plane of a microscope objective. The light from the pinhole is focused by the objective and a demagnified image is created. The optics are mounted on a precision motorized stage to allow the spot to be moved around

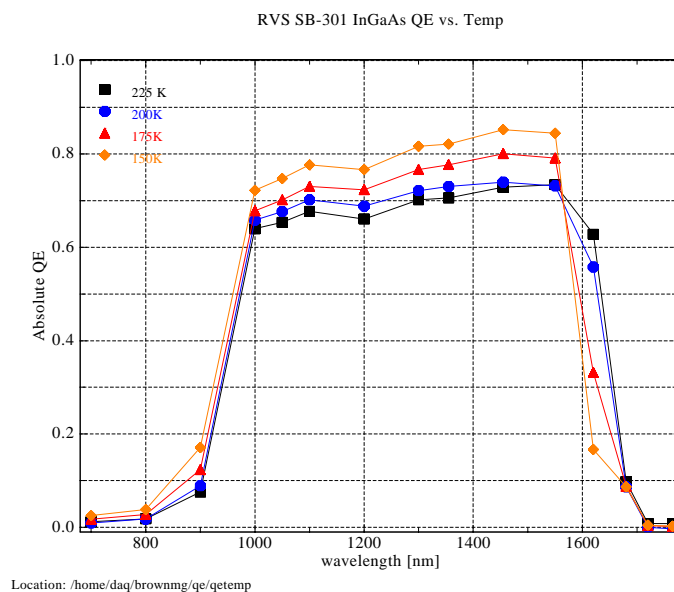


Figure 1.1: Quantum efficiency as a function of wavelength for several different operating temperatures for the Raytheon Vision Systems (RVS) SB-301 InGaAs detector.

on the device. If the spot is small enough, it can be used to probe small variations in the response of the device, as well as large scale variations.

This is a common solution for probing semiconductors, but it is usually done for CCDs [4] and so uses visible light. A pinhole projector for NIR light has never been previously created, nor had one for visible light been created that had the precision required to perform these measurements. It was therefore necessary to design such a system for the NIR, build it, and test it in-house. The University of Michigan now has such a system, fondly named “Spot-O-Matic.”

This thesis documents the design, creation, and testing of the Spot-O-Matic. It also presents results obtained from using the Spot-O-Matic to fully characterize the Raytheon Vision Systems SB-301 NIR InGaAs detector. This InGaAs detector is similar to those that will be used on the SNAP focal plane array. It has slightly larger pixels ($20\mu m$ as opposed to $18\mu m$ for the HgCdTel detectors), and its quantum efficiency cuts off at about $1600nm$ (see Fig. 1.1). However, it still serves as an excellent test device to prove the functionality of the Spot-O-Matic, as well as provide valuable insight into testing procedure and NIR detector design.

Chapter 2

Experimental Setup

2.1 Spot-O-Matic Optics

The concept of a pinhole projector is a common solution for probing the response of semi-conductors. However, the fact that the SNAP devices must be probed in the NIR, together with the fact that we require uncommonly precise measurements, required that a unique experimental setup be designed.

The SNAP photodetectors must be tested in the same conditions in which they will exist on the actual telescope. Therefore, when performing any measurements, the detectors reside in a cryogenic dewar, keeping them in a high vacuum and at 140K. The only light allowed to strike the detector must pass through a window on the front of the dewar. Due to mechanical restrictions, the front of this window is about $3cm$ away from the front of the detector. Therefore, any objective used must have a working distance greater than $3cm$.

In addition to this focal length requirement, we also require a lens system that is suited for both visible and NIR light, since we do not wish to limit the functionality of our apparatus in the visible spectrum. Therefore, it must be chromatically corrected throughout the entire range from the visible up to $1700nm$. In addition, we wish to project the smallest spot possible, which means choosing a lens with a large numerical aperture (NA), since this would give us a lower diffraction limit for spot size according to

$$RP = \frac{0.61\lambda}{NA}, \quad (2.1)$$

where RP stands for resolving power. Taking all this into account, I selected the Mitutoyo M Plan 10x microscope objective as the primary lens. It has a working distance of $30.5mm$, demagnification of 10x, and a numerical aperture of 0.26, which according to the above equation, with $1700nm$ light, would produce an Airy disk with a distance from peak to first minimum of $3.99\mu m$. Of course, any projected spot will be larger than this, since it will be the convolution of the demagnified image of the spot with the Airy disk function, but anything around this size will be suitable for the measurements we intend to make.

The microscope objective is infinity corrected, meaning it requires an auxiliary tube lens to function correctly. Therefore, the InfiniTube NIR chromatically corrected tube lens from Infinity Optics was selected. A tube lens is essentially a lens that is designed to couple with the objective, but which has a much longer focal length than the objective. The Mitutoyo

objective has an effective focal length (not to be confused with working distance) of 20mm , while the tube lens has a focal length of 200mm . They are coupled so that near-parallel light exits the tube lens and enters the objective. Since the focal length of the tube lens is so much longer than that of the objective, this allows one to adjust the distance between the pinhole and the tube lens (and thus the magnification) without altering the working distance much at all. The combination can be thought of as sort of an “optical lever.” Fig. 2.1 shows the experimentally measured effect of pinhole position on magnification and working distance. The tube lens-objective combination allows us to move the pinhole to obtain a desired spot size without altering the crucial working distance.

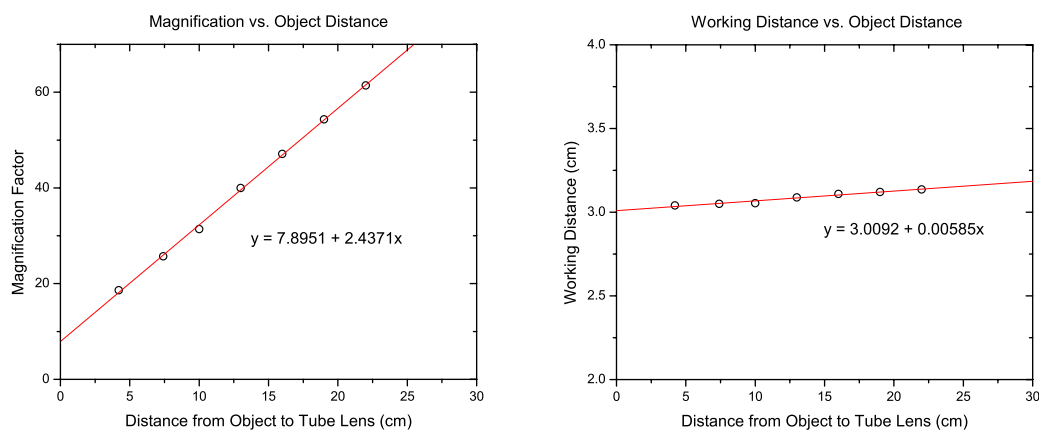


Figure 2.1: Measured magnification and working distance as functions of the distance between the pinhole and the tube lens.

Since the actual size of the projected spot will be the convolution of the Airy disk function of the tube lens/objective and the demagnified image of the pinhole, the smallest pinhole possible should be chosen to minimize spot size. However, the smaller the pinhole, the higher the error on roundness. Therefore, a $10\mu\text{m}$ diameter pinhole was chosen as suitable, since with even 10x demagnification, it should not add much size to the diffraction limited spot.

A $10\mu\text{m}$ pinhole was also chosen because, to be precise, the resulting spot on the detector will involve both the the Airy disk function of the tube lens/objective as well as the Airy disk function of the pinhole. The smaller the pinhole, the wider the central maximum of the resultant Airy disk will be. We want this maximum to be wide enough so that is completely fills the tube lens, so that we will not have rings in our projected spot.

The optics of the Spot-O-Matic for the current configuration are shown in Fig. 2.2. Light travels from our light source, through a liquid light guide, and into the optical assembly. Upon leaving the light guide, the light is focused by a convex lens and then split by the beam splitter. The weaker beam goes to the feedback regulating diode. The stronger beam passes through an interchangeable narrow bandpass filter to select wavelength, and then through a diffuser that sits directly before the pinhole. The diffuser ensures that the pinhole is uniformly illuminated, which makes for a uniform projected spot. The pinhole resides at

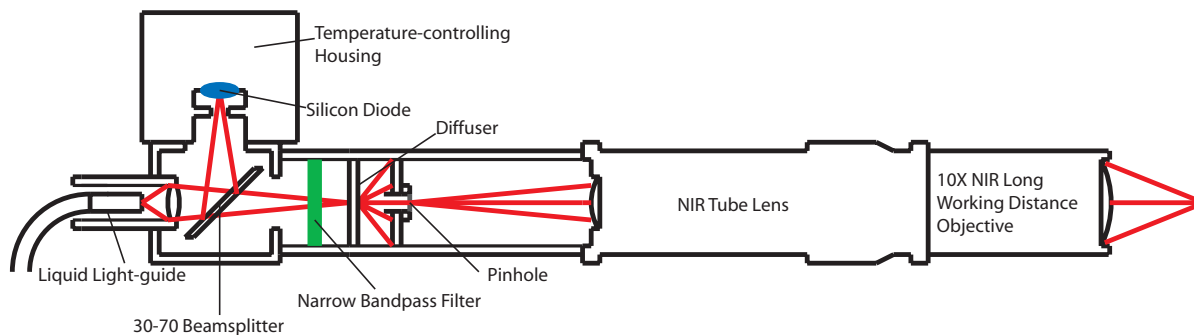


Figure 2.2: Diagram of the optical components of the Spot-O-Matic.

the focal plane of the objective/tube lens combination. Light from the pinhole enters the tube lens and objective, and a demagnified image is produced on the device. Note that the demagnification depends on the distance from the tube lens to the pinhole. All of these optical components are held together inside of an interchangeable tube assembly purchased from Edmund Industrial Optics. These tube assembly pieces are 1 inch in diameter, and are anodized black and threaded on the interior to reduce stray light. Even with this baffling, it was still necessary to line the inside of the tube between the pinhole and tube lens with dull black paper, in order to further reduce scattered light.

The light source had to be specially selected to meet the precision required in our measurements. Since we are attempting to verify that the response of the device is uniform to within 2%, we must have a spot with ultra-constant intensity. Therefore, the Photomax light source from Oriel was selected for its superior stability and the option to include feedback regulation. Fig. 2.3 shows the intensity over time of the Photomax light source with and without feedback regulation, according to the manufacturer [5]. It uses a 50W quartz tungsten halogen (QTH) lamp, focused by an ellipsoidal reflector to maximize light output at the other end of the liquid light guide. It is essential that the light coupling be excellent, since we are using filters with bandpass of only $10nm$ FWHM, so we need all the light we can get. A liquid light guide was selected over standard fiber optics since the transmission of a liquid light guide is nearly constant over the entire visible and NIR spectrum.

The power supply associated with the Photomax light source produces a very stable current, with load regulation of $\pm 0.05\%$ [6]. However, even with an ultra-stable current, the light output of the lamp can vary widely, due to lamp age, filament erosion, gas adsorption or desorption, and particularly ambient temperature. Because of these factors, the actual light output can vary by as much as 4% over the span of 24 hours. Since pixel scans often last that long, this is not stable enough for our purposes. Therefore, we purchased the optional Light Intensity Control System from Oriel to further stabilize the output. As described above, the output is feedback regulated by light directed by the beamsplitter into a temperature-controlled diode. Using feedback regulation, the output is stable to much better than 1% over more than 100 hours.

Since we are attempting to probe structure in the PRF at the micron level, we must be able to focus the spot on the device and move it around with very high precision. Therefore,

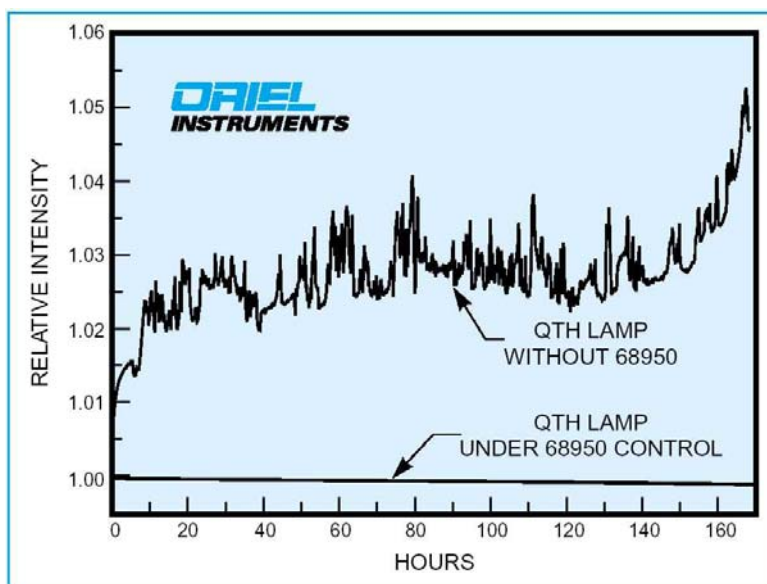


Figure 2.3: Sample light intensity over time of Photomax light source, with and without feedback regulation. Figure courtesy of Oriel Instruments. [5]

I selected the MM-4M-EX 3-axis computer controlled linear translation stage from National Aperture. The stage has step size of $0.0749\mu\text{m}$, is accurate to $1.0\mu\text{m}$ per inch of linear travel, and has repeatability of $\pm 0.5\mu\text{m}$. This is accurate and precise enough to allow us to probe pixel response at the micron level.

In order to control the motion of the stage, extensive software had to be written to control motion specific to our measurements. This software was written entirely by Michael Borysow, and without it none of these measurements would be possible. [7]

One problem we encountered with the stage is that the position of the stage is monitored by rotary encoders on the motors. Therefore, any change in direction will introduce a new error in position due to slider backlash. This backlash offset is significant and not constant. Thus, whenever one of the axes switches directions, the spot's actual physical position and the position recorded by the encoder can be off by as many as several microns. This backlash has been thoroughly characterized by Michael Borysow [7], and his results are presented in Fig. 2.4.

In theory, when moving the spot parallel to the plane of the device, this offset can be corrected by always moving in the same direction, and if a reverse is necessary, by overshooting and coming back from the original direction. Therefore, the backlash offset will have been introduced twice, in opposite directions, and should effectively cancel out. Unfortunately, this is not always the case, since the backlash is not always constant, nor symmetric. But fortunately, the boundaries of pixels can be used as reference points to correct for any error.

However, when focusing the spot on the device, that is, moving toward and away from the device, it is necessary to have an absolute knowledge of where optimal focus is, since it is sometimes necessary to move to a different focus position, and then return to the original position, without having to account for the direction of approach. This requirement led us

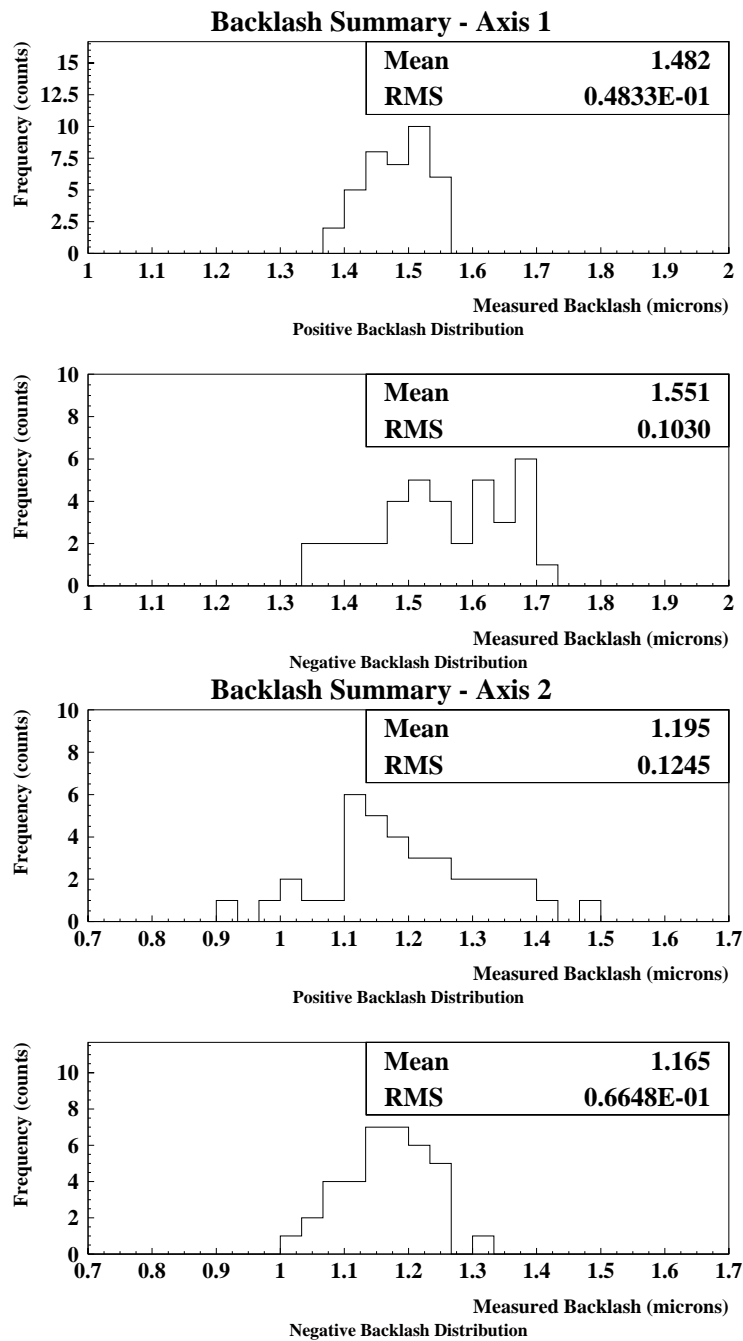


Figure 2.4: Histograms of backlash results for the x and y axes of our stage. It can be seen that not only is the backlash as large as several microns at times, but it is also not constant, making it difficult to compensate for it in software. Graphs courtesy of Michael Borysow [7].

to install a linear encoder on the z -axis slider. With a resolution of $20nm$, it not only gives us a much more accurate knowledge of our position, but it also eliminates the effect of slider backlash. Therefore, once the focus position is established, it will always be at the exact same encoder reading, as long as everything is firmly mounted.

2.2 Experimental Setup

The Spot-O-Matic, with the stage, optics, and linear encoder all together, is shown in Fig. 2.5. Refer to Fig. 2.2 to compare with the inner-workings of the optics.




Figure 2.5: The fully functional Spot-O-Matic, positioned in front of the dewar window, ready to perform measurements on the InGaAs detector.

The Spot-O-Matic is securely mounted on an optical bench to prevent everything from moving and to isolate it from external vibrations. It sits inside a dark box to prevent stray light. A data acquisition computer is responsible for controlling the motion of the stage as well as taking exposures from the device.

Chapter 3

Experimental Procedure

3.1 Characterizing the Spot Size

Before any measurements can be performed, it is crucial to know exactly how large of a spot we are projecting on the device. We can estimate its size based on our knowledge of the numerical aperture and the demagnification of the objective. But this will not tell us if there are any unexpected features of the spot due to particulates on the lenses, or defects in the pinhole, for example. We can expect something resembling an Airy disk, since with a $10\mu m$ pinhole, the spot size is dominated by diffraction, but it is crucial that we be sure.

In order to accurately determine the spot profile, we use a procedure called the knife-edge scan, as illustrated in Fig. 3.1.




Figure 3.1: Diagram of knife-edge scanning procedure.

First, the spot is roughly focused on the edge of the razor blade. The Spot-O-Matic is then scanned perpendicularly to the edge of the blade, thereby obscuring the spot incrementally. The intensity of the portion of the spot reaching the detector is recorded versus the position relative to the edge of the razor blade. Fig. 3.2 shows the results from knife-edge scans performed with $1050nm$ and $1550nm$ light at focus.

The top two panels show the intensity on the detector versus position. This measured intensity is actually the integral of the one-dimensional profile of the beam intensity. Thus, by taking the derivative of the top panels, we can determine the spot intensity profile. In order to take the derivative of discrete data points, we smooth the initial data with a spline

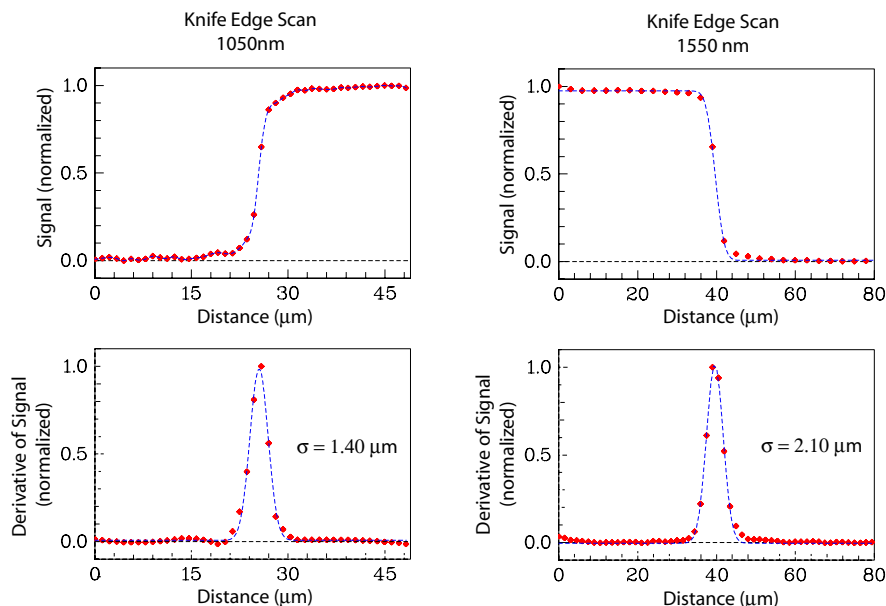


Figure 3.2: Knife-edge scans for 1050nm and 1550nm light at focus. The top panels show intensity recorded on the detector as a function of spot position relative to the edge of the razor blade. The bottom two panels show the derivative of this intensity plot, which gives us the intensity profile of the spot.

function, and then take the derivative. These discrete derivative values are plotted in red in the bottom panels. Since we expect our diffraction-limited spot to be very close to an Airy disk, and because the central maximum of an Airy disk function can be approximated quite well with a Gaussian function, we fit a Gaussian to these data points. This Gaussian is plotted as the dotted blue line in the bottom two panels of Fig. 3.2. The Gaussian does indeed fit the data very well. Now that we have a simple function to work with, we extract σ from the Gaussian and use that as our measure of spot size.

In order to obtain the profile of the *focused* spot, we have to be sure that the spot is exactly focused on the edge of the razor blade. This requires us to do many knife-edge scans, each time moving the focus forward slightly, until we find the smallest spot. We extract σ for each knife-edge scan, and then plot σ versus z -axis position. When we do, we get a plot like that in Figure 3.3.

In the far field, spot size should change linearly with distance from focus. Near the focus, spot size will change more like a quadratic function. Therefore I fit a linear regression to the points farthest from the anticipated focus, and I fit a second-order polynomial to the points near focus. The results are nearly perfect for this data. The linear fits both point to a focus point between the lowest two data points, as does the quadratic. The curve is quite flat within several microns of this expected focus point, which is consistent with the manufacturer's specification that the microscope objective has a depth of focus of $4\mu\text{m}$. Our lowest data point is very close to the minimum of this curve, so we can safely use this point as our minimum spot size.

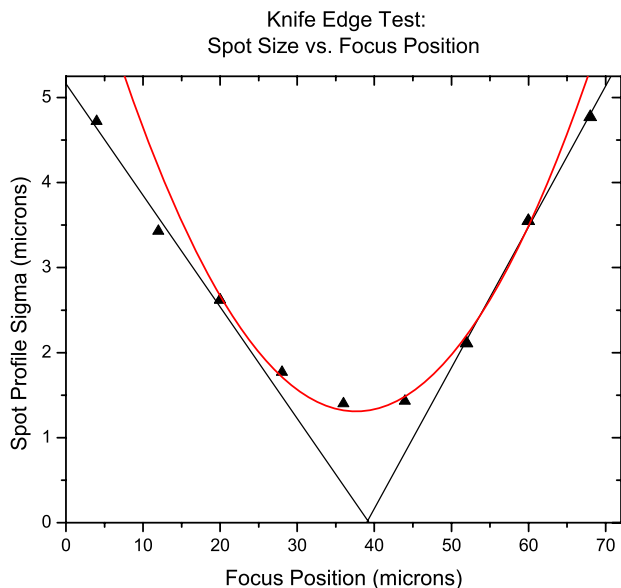


Figure 3.3: A plot of σ versus focus position for $1050nm$ light. The black lines represent the variation of spot size with distance far from focus. Near the focus, the spot size changes more like a quadratic function, shown with the red curve. The minimum σ is $1.40\mu m$. The profile of this spot is plotted in Fig. 3.2.

We find that for $1550nm$ light, we get a spot with $\sigma = 2.10\mu m$, and $\sigma = 1.40\mu m$ for $1050nm$ light. This agrees with the wavelength dependence of the diffraction limited spot size, since

$$\left(\frac{1050nm}{1550nm}\right) \cdot 2.10\mu m = 1.42\mu m.$$

Now that we have an accurate knowledge of what size spots we are working with, we can proceed to conduct scans on the device.

3.2 Focusing the Spot on the Device

Since we know the size of our spot at focus, in order to produce accurate measurements, we must ensure that the spot is focused on the device when we do any scans. This requires use of a procedure called the virtual knife-edge scan.

The virtual knife-edge scan functions in much the same way as the real knife-edge scan, except that now, instead of a razor blade, the edge of a pixel is used as a “virtual knife-edge.” The edge of a pixel is by no means a sharp cut-off (due to precisely the things we are trying to measure – intra-pixel variation, lateral charge diffusion, etc.) but it should still allow us to find when the spot is focused on the device.

As with the real knife-edge scan, we will scan across the edge of the pixel many times, each time moving slightly in the z -axis. We plot the intensity recorded in the pixel versus position, and we obtain graphs that look like Fig. 3.4.

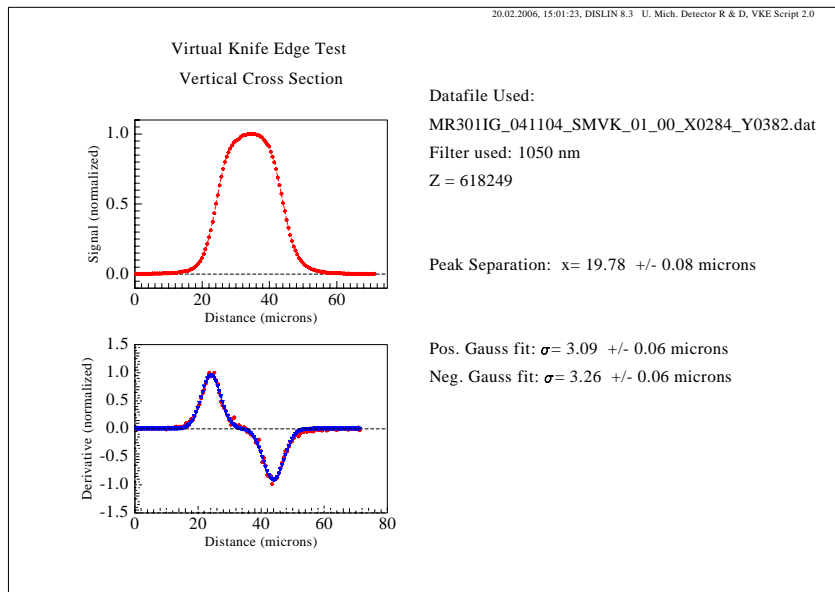


Figure 3.4: A virtual knife-edge scan through a sample pixel at optimal focus with 1550nm light.

Fig. 3.4 is basically a one-dimensional (1-d) slice of the PRF, convolved with the intensity profile of the spot. Therefore, theoretically, the smaller the spot, the sharper the slopes will be on both sides. In order to determine how sharply the intensity rises on both sides, we again smooth the raw data with a spline and take the derivative. We then fit each peak to a Gaussian, and extract σ . We assume that the smallest σ will result from the smallest spot, and thus the position when the spot is in focus. We can use either the rising edge or the falling edge to find focus. To be consistent, we plot the Gaussian width of the left edge as a function of z -axis position, and again, we fit a second order polynomial to the inside data and a line to the outer data points. Fig. 3.5 shows an example of this. The lines and the parabola do not agree as well as they did for the real knife-edge scan. Nevertheless, we can still get a confident measurement of when the spot is focused on the device from the second-order polynomial fit.

Once this is accomplished, we know permanently the focus position for this region of the device, thanks to the linear encoder on the z -axis. It is important to note that this focus position applies only for this place on the device, and a small surrounding region. This is because the x - y plane of stage movement is not in general parallel to the plane of the surface of the device. Also, it is not guaranteed that the device is completely flat. Therefore, if we move the stage in x or y , the spot may no longer be exactly focused on the device. This is not a problem as long as we stick to a small region on the device, say, perhaps a 10×10 region of pixels, since the optics have a depth of focus of $4\mu\text{m}$, giving us some leeway.

In order to scan larger regions of the device, we must create a focus map of the device.

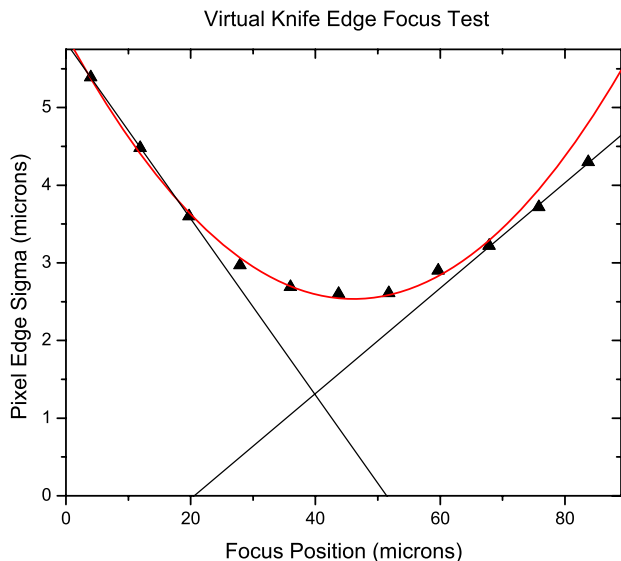


Figure 3.5: A plot of σ versus z -axis position from virtual knife-edge scans performed on a sample pixel at $1550nm$.

This requires us to perform virtual knife-edge scans at various regions of the device and find the focus at each point. We can then determine the plane (if it is a plane) of the device relative to the x - y plane of the stage, and correct in z for any movement in x or y , and therefore keep the spot focused on the device.

We have not accomplished a full focus map yet, since it requires rather complicated software additions in order to automatically correct in z for small steps in x and y . Eventually this will have to be implemented when we test actual SNAP devices.

3.3 Tests to be Performed

This thesis documents the results of many tests performed on the RVS InGaAs test device. The following are descriptions of how these tests are performed. The following procedures will not change when conducting tests on HgCdTe devices.

Much has already been said about the real knife-edge scan for characterizing spot size. It is important to add that the above description only detailed how to obtain a 1-d intensity profile of the spot. For the measurements in this thesis, we have only used the results from a 1-d scan, and so we are assuming that our spot is circularly symmetric, and that if a knife-edge scan were performed in the perpendicular direction, it would give the same results. This assumption is acceptable, as long as the spots are small. If the spots are indeed small, they will be dominated by diffraction and any aberrations that would make the spot non-symmetric would be blurred out. Also, any deviations from circular would not make much difference since the size of the spot is so much smaller than the pixels we are scanning.

For our $1.40\mu m$ and $2.10\mu m$ spots, this is the case.

However, it is sometimes necessary to use larger spots to perform a variety of scans. In these situations we use a larger pinhole in the Spot-O-Matic and produce a bigger spot. For these larger spots (about $\sigma = 7\mu m$) diffraction no longer dominates the geometric optics, and so any defects in the pinhole or aberrations in the optics could produce deviations from circular symmetry. In this case, it would be necessary to do two perpendicular knife-edge scans, and then approximate the spot as an ellipse. One such set of large-spot data is presented in Chapter 5, but since this is a test device, the circular approximation was deemed acceptable. For more precise measurements in the future, more accurate characterization of the spot may become necessary.

The virtual knife-edge technique has already been described, and it has been discussed how it is used to focus the spot on the device. This scan is also useful for extracting much information about the pixel response function.

The idea of the virtual knife-edge technique is to scan the spot directly through the middle of the pixel, perpendicular to the edges of the pixel. From this scan, we can determine the width of the pixel, the lateral charge diffusion, any gaps between pixels, and several other things. However, the 1-d virtual knife-edge technique is limited because we can never be exactly sure that we are scanning through the center of a pixel, and as will be seen in Chapter 4, this will affect measurements. For this reason, in order to accurately extract the PRF, we use a full pixel scan.

A pixel scan is essentially a 2-d virtual knife-edge scan. The spot is scanned across the pixel in x with steps of 30 counts. Then, the spot is brought back to the starting position, moved up by 30 counts in y , and then the scan in x is repeated. This scanning procedure is illustrated in Fig. 3.6. We thus obtain a 2-d mesh of pixel response as a function of spot position. Note that images are also taken when the spot is located *outside* the physical area of the pixel, since we recall from Chapter 1 that the PRF is non-zero outside the physical area of the pixel.




Figure 3.6: Schematic view of a pixel scan. The red line shows the motion of the spot, and the dots represent places where the stage stops and an image is taken.

That is the procedure for determining the PRF of a single pixel. In order to see if there are any dips between pixels, or if the device is flat to within 2%, we must sum the PRFs

of adjacent pixels. Since a pixel's PRF is larger than the physical pixel size, the PRFs of adjacent pixels will overlap. Thus a photon striking in the region of the overlap might be recorded as a count by either pixel. So to determine the overall QE of a single point on the device, we must sum the response of *all* pixels with a non-zero PRF at this point.

Usually a 5×5 region of pixels is selected, and a pixel scan is performed as described above over the whole region. The difference is that when scanning a single pixel, at each scan position we simply plot the readout from the pixel of interest. When scanning a region, at each scan position we take the readout from *every* pixel in the region and sum them up. Those pixels in the region with overlapping PRFs at that particular scan position will have their responses added together to reveal the total number of photons collected at that point. Those pixels with PRFs that do not extend to that particular scan position will add nothing but noise to the sum.

Theoretically, we should get a flat graph. This is because no matter where the spot is on the device, the photons are being collected somewhere, and so when all the pixels are summed up, we should get the same response no matter where the spot is shining. If we see something other than a flat graph, this means that photons are being lost when the spot is shining on certain regions of the device, or that the device has a non-linear response. (It turns out that the InGaAs device *does* have a non-linear response, the implications of which will be discussed in Chapter 5.) For the time being, a 5×5 region is about as large as we can scan, since we do not have a focus map for the device, as was mentioned above. To ensure that we are staying in focus, we restrict ourselves to a small region.

Chapter 4

Characterizing the Pixel Response Function

4.1 Discussion of the Pixel Response Function

Something has already been said about the importance of knowing the pixel response function. As discussed in Section 1.2, knowing the PRF will allow us to determine if the device is flat everywhere to within 2%, and also to reconstruct dithered astronomical images, thereby recovering some of the information lost to undersampling.

As was mentioned in Chapter 1, the measured PRFs on page 21 deviate from the ideal 2-d square function for several reasons. The first reason is possible defects in the pixel. These could be caused by any number of factors in the manufacturing process, and range from the rather mundane “pock mark” of reduced sensitivity to the considerably more bizarre (see Chapter 7).

The second reason is lateral charge diffusion. At this point, a more thorough discussion of lateral charge diffusion is in order. Normally when a photon strikes a pixel, it is absorbed at some depth within the substrate, freeing an electron, which is then accelerated toward the junction by the bias voltage and collected on a capacitor to be read-out by the electronics. Occasionally, an electron that was freed inside a certain pixel will travel into a neighboring pixel and be collected by the capacitor for that pixel. Therefore, even though no light is striking a pixel, it can still have a non-zero read-out. Obviously, an electron is more likely to travel to a neighboring pixel if it is freed near the edge of a pixel rather than in the center. Therefore, instead of a 2-d square function, the PRF will be highest in the center, since most electrons freed here will be collected by that same pixel, but as you shine light closer to the edges, the response will be lower since some of the electrons are escaping. In addition, the PRF will be non-zero outside of the physical size of the pixel. That is, even if light is not shining on the pixel itself, it will obtain some electrons from a neighboring pixel that is being illuminated. These factors will cause the PRF to look not like a square function, but rather like a 2-d Gaussian with a flat top.

Complicating the PRF further is the fact that it will be dependent on the wavelength of light striking it. This is because photons with longer wavelengths travel further into the substrate on average before being absorbed. Therefore, once an electron is freed, it tends

to be closer to the back of the substrate and to the junction, where the electric field is stronger, and therefore has less chance to escape into a neighboring pixel. Thus, lateral charge diffusion will be greater for shorter wavelengths of light. This means that the PRF will become even less square-like. The flat top on the Gaussian will tend to disappear toward shorter wavelengths as pixels extend their “reach” into their neighbors.

To demonstrate the problem of lateral charge diffusion, I will show a simple example. Figure 4.1 shows the response of several pixels. Our small spot is shined in the center of one of the pixels. Since our spots are much smaller than the pixel size ($1.40\mu\text{m}$ for 1050nm , $2.10\mu\text{m}$ for 1550nm), we can be confident that almost no light is shining directly onto any pixels except the brightest one. However, it is clear that the neighboring pixels are reading out well above the noise level. Notice also that the surrounding pixels are responding relatively higher for the 1050nm spot than for the 1550nm spot. This is *not* due to the fact that the 1550nm spot is larger than the 1050nm spot, since both spot sizes are very small compared to the size of the pixel. The reason as described above is that 1050nm photons penetrate less deeply into the substrate, and so lateral charge diffusion is greater.

There is a third, purely experimental reason why the measured PRFs on page 21 will deviate from ideal square functions. The reason is that the plots on page 21 are not actually the *real* PRFs of the device. Theoretically, in order to really obtain the pixel response at a specific *point* on the device, we would have to scan it with an infinitesimally small spot, i.e. a delta function. Our spot is small, but it is still not small enough to consider it a point. Therefore, what we will obtain from our pixel scan will be a convolution of the 2-d intensity profile of the spot with the *real* PRF of the pixel. Therefore, in order to extract the real PRF of the pixel, we need to deconvolve the data we get from the scan with the intensity profile of the spot. So, even if the real PRF were a 2-d square function (which it’s not), we still would obtain something that looks like a square function with sloped edges. After taking lateral charge diffusion into account, this unavoidable experimental fact serves to make the observed PRF even more Gaussian-like.

4.2 Extraction of the Pixel Response Function

Finally it is time to actually map the 2-d pixel response function. To do so, we return to our small spots of $1.40\mu\text{m}$ and $2.10\mu\text{m}$. We perform two scans of the same pixel as described in Sec. 3.3, one at 1050nm and one at 1550nm . The results are shown in Fig. 4.2.

Several things are immediately apparent. First, the pixel response seems to be very smooth over the surface of the pixel. There are not dips or small-scale structure that would suggest that charge is being lost at any point inside the pixel.

Second, we notice that the effects of lateral charge diffusion are significant. The plots resemble a Gaussian much more than they do a square-function. We also notice that the PRF is larger than the physical size of the pixel. This is also what we expect due to lateral charge diffusion.

Third, we notice that the results for the two wavelengths look slightly different, as we expect they should. The difference might not be immediately obvious, but the 1550nm scan looks slightly more square and has a flatter top. This is also what we expect, since charge diffusion is greater for 1050nm .

It is important to keep in mind that what we are looking at are convolutions of the PRF with the 2-d intensity profiles of the spots used in the scans. Ultimately, we would like to see the deconvolved PRF for both wavelengths in order to compare. Since the spot sizes used in the scans are different for each wavelength, we have introduced unequal experimental error into the two data sets. It is important that we recover the real PRF at both wavelengths for use in reconstructing dithered images. However, it is not crucial for verification of the photometry requirements, as will be described in Sec. 5.4.

However, 2-d deconvolution is not an easy task. According to the Nyquist Theorem, to avoid significant information loss, the highest frequency component of the spot intensity profile can be no higher than half of the sampling frequency of the convolved PRF. [8] Unfortunately, for these scans, images are taken every 30 encoder counts (see Sec. 3.3), which means the data points are spaced $2.25\mu m$ apart. Since the largest spot we use has $\sigma = 2.10\mu m$, we obviously need finer sampling of the convolved PRF to deconvolve these scans. We could repeat the scans with smaller steps, but these scans take a very long time as it is, often as long as 10 hours. Halving the step size would quadruple the scan time, and then we risk bringing the long-term stability of the light source into the equation. Luckily, there is a simpler way to extract the important data.

We take advantage of the fact that the response seems to be very symmetric with respect to x and y . Therefore, we can simply do a virtual knife-edge scan through the center of the pixel, sampling more often, and then deconvolve this with the 1-d intensity profile of the spot.

Therefore, we do a virtual knife-edge scan through the center of the pixel at $1050nm$ and $1550nm$, with step size of 7 encoder counts, or $0.52\mu m$. We make sure that we are going through the center of the pixel by taking advantage of lateral charge diffusion. We make sure that the pixels above and below the pixel of interest read nearly the same while the spot is traveling through the center pixel. This is crucial, because if we do not scan through the exact center of the pixel, the resulting profile will be different, as one can see from Fig. 4.2. The results are shown in Fig. 4.3. Here the difference between the two wavelengths is even more apparent.

Once we deconvolve this data, we will have completely determined the PRF of this device. We will have all of the information necessary to reconstruct dithered images, and part 1 of our task will be completed. Once completed, we can then move on to summing the PRFs to verify the photometry requirement. However, deconvolution has proved extremely difficult. High frequency components in the convolved signal cause serious problems when taking the Fourier transform. Better ways to deconvolve are being looked into by the SNAP group, so I leave this task beyond the scope of this thesis. Even though we have not obtained the true PRF of this device, we can still move on to verifying the photometry requirement. The raw, convolved data turns out to be just fine for this task, as will be seen.





Figure 4.1: The top two panels are raw FITS files of a small section of pixels illuminated at each wavelength, viewed with the SAOImage DS9 software. The difference in the response of the neighboring pixels is slight, but noticeable. The bottom panels are histograms of the pixel read-out, normalized, and with the range adjusted to highlight the difference in lateral charge diffusion. The diffusion is slightly higher for $1050nm$.

Single Pixel Response vs. Spot Position
X228_Y353
1050nm

Maximum Flux: 10931.0 ADU



Single Pixel Response vs. Spot Position
X228_Y353
1550nm

Maximum Flux: 10634.0 ADU




Figure 4.2: Single pixel scans at 1050nm and 1550nm. The grid on the bottom represents the physical size of the pixel ($20\mu\text{m}$).

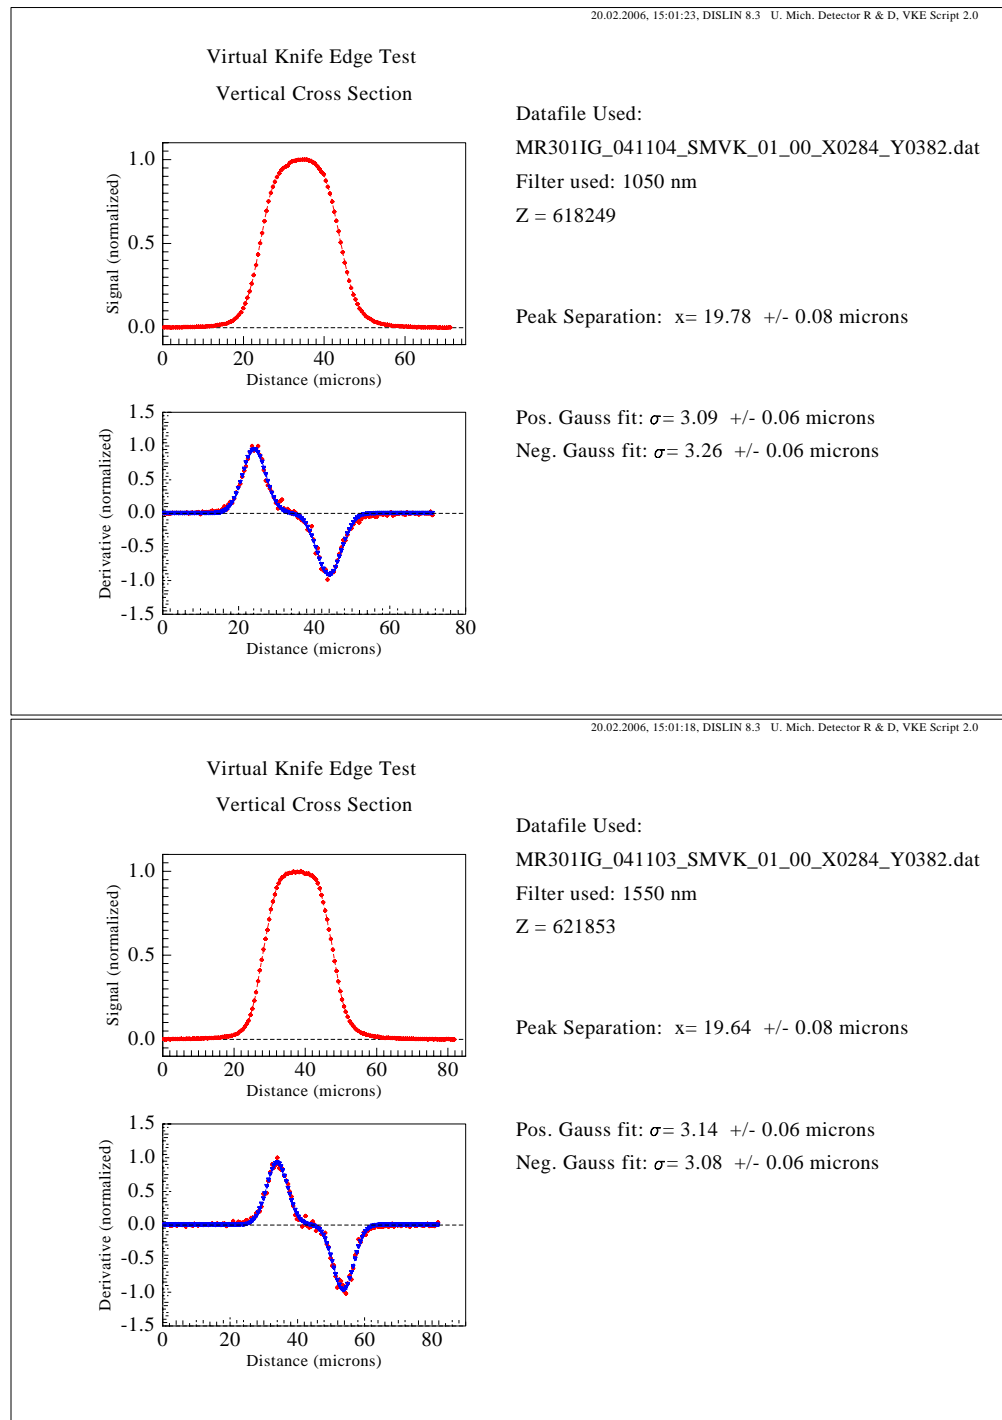


Figure 4.3: One-dimensional virtual knife-edge scans through the center of a pixel, at 1050nm and 1550nm. The derivatives of the data points are shown below, with the Gaussian fit of the peaks in blue.

Chapter 5

Photometry Requirements

Part 2 of our task is to verify the SNAP requirement that the response of the device be flat to within 2%. As mentioned in Chapter 1, there are essentially three possibilities for deviation from flatness: variation within a pixel, dips between pixels, or differences in average sensitivity among individual pixels. The last one is not so much a concern as this can be rather easily fixed with a simple flat-field correction. Also, we have seen nothing in the convolved PRFs to suggest that variation within a pixel will pose a problem. Both PRFs appeared smooth with no dips or pock marks (you may flip ahead to Fig. 7.1 to compare with a PRF in which this *is* a problem). Therefore, if we are going to see any deviation from flatness, we expect that it will be found as sensitivity dips between the pixels. In order to see this, we must sum adjacent PRFs. Even though we have not extracted the true PRF, for the purpose of verifying the photometry requirement, we can proceed with the raw, convolved data as though it were the real, deconvolved PRF. It will not affect the results. Sec. 5.4 will explain why this is the case.

5.1 Initial Search for Dips

As previously mentioned, we expect that there should be regions between pixels where QE is lower. We expect this because the junctions on the InGaAs device are located below the center of the pixel. Therefore, we expect that electrons freed near the edges or corners of pixels will be collected with less efficiency, since they are a maximum distance away from any one junction.

We could perform a scan over a region of multiple pixels as described in Sec. 3.3. However, there is a quicker and easier way to see if we are correct. As was done in Sec. 4.2 in order to deconvolve the PRF, we simply do a virtual knife-edge scan in lieu of a full region scan. We do a virtual knife-edge straight through the center of three adjacent pixels, and plot the readout from each pixel on the same set of axes. We then sum the readouts from all three pixels at each point in the scan. The results are shown in Figure 5.1.

We do not see what we expected. We do not see any dips between pixels in the 1050nm scan. In the 1550nm scan, we actually see small regions of *increased* sensitivity. This is very counterintuitive. Seeing no dips at all is within reason, but seeing bumps rather than dips makes no sense at all. The answer lies in the the fact that the InGaAs detector exhibits a

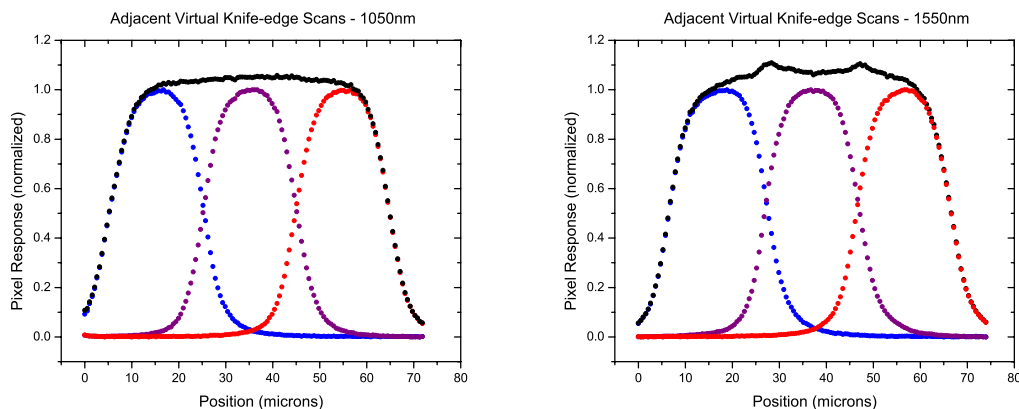


Figure 5.1: Virtual knife-edge profiles for 3 adjacent pixels at both wavelengths. The response of all 3 is summed at each point and plotted above the individual profiles.

non-linear response.

5.2 Non-linear Response

The reason for this apparent increased sensitivity of the device along pixel boundaries is that this particular device has a non-linear response. This means that there is not a constant relationship between read-out counts and photons collected. In fact, the pixels of the InGaAs device respond proportionately higher when the number of photons collected is relatively low. Put another way, doubling the number of photons collected by a certain pixel will *less than* double the read-out. Therefore, we see the highest response near the corners of pixels, since when the spot is at this point, its intensity is divided among four pixels. When the response of these four pixels are summed up, the result will be higher than when the full intensity of the spot is contained within a single pixel. Therefore, these areas of “increased sensitivity” are not variations in the QE of the device, but rather a characteristic of the electronics. Because of this, we cannot determine if there is any reduced QE at the pixel boundaries, since it is masked by this non-linearity.

Since the non-linearity of the device should be measurable, the part of the variation due to this effect should be correctable, both in the lab and on the SNAP telescope. Therefore we would like to correct for this feature of the electronics so that we can measure the true flatness of the device. Once we do, we expect that the anticipated dips between pixels would become visible, since we still expect to see these dips based on the design of this device.

In order to perform this correction, we would need to shine the spot in the center of a pixel, and then increase the intensity incrementally while recording the read-out, thus mapping the response curve of the device. The response could thus be translated into actual physical light intensity on the device, and non-linearity would no longer be a problem. At the time of writing this, this has not been carried out for the InGaAs device, since it would be time-consuming and this is only a test device. However, the Spot-O-Matic is quite capable of

performing this test should it become necessary in the future. For the measurements in this thesis, we accept this non-linearity and its associated deviations from flatness as unavoidable features of this device.

A brief discussion is in order as to why we only see the bumps on the $1550nm$ scan. There are two possible reasons. First, the non-linearity of the device might be wavelength dependent. Second, the greater lateral charge diffusion at $1050nm$ may have “smeared out” these bumps so that they are no longer visible. The second option seems most likely. From Fig. 4.3 we see that lateral charge diffusion at $1050nm$ tends to round off the relatively sharper corners seen on the flat top at $1550nm$. It is then reasonable to assume that it has an analogous effect on these inter-pixel bumps.

5.3 Multiple Pixel Scans

In order to more fully understand this response between pixels, and to verify that this device is flat to within 2%, we perform a 3×3 pixel scan as described in Section 3.3. As usual, we do one at $1050nm$ and one at $1550nm$. The results are shown in Figure 5.2.

Fig. 5.2 shows the response of 9 pixels at $1050nm$. The response is quite flat, as expected from Fig. 5.1. The response has a Gaussian standard deviation of only 0.0085, as shown in the histogram. This is excellent flatness, well within the acceptable 2% range for the SNAP telescope.

Fig. 5.3 shows the response of the same 9 pixels at $1550nm$. Here, we see the ridges between pixels, as we expected from the two dimensional scans. However, even with this potentially correctable error, we obtain a standard deviation about the mean of only 0.0141, as shown in the histogram. Again, this is well within SNAP specifications.

These 3×3 scans are actually subregions of larger scans that were performed. The larger scans did not look as good as these scans, and had larger variations when the extended regions were included, for reasons that are discussed in Chapter 6. However, this was due to light intensity drift in the light source. Indeed, proving the flatness of a 3×3 area is all that is required to prove that the device is within SNAP specifications, as long as the pixels in the region are typical of the rest of the device. A histogram of a 3×3 scan takes into account the deviation from flatness due to non-uniformity within a pixel as well as any sensitivity decrease (or increase) at all 4 boundaries of a pixel. Extending the region of the scan would only introduce redundant variation, which would not affect the standard deviation about the mean. Any variations on larger scales can be corrected with a flat field, and so would not hinder SNAP specifications.

We have thus proven that the RVS InGaAs detector would pass SNAP specifications for flatness. We have also determined the PRF of this device, minus deconvolution. These were essentially the two main goals that Spot-O-Matic was commissioned to fulfill. With these results we are technically done. The Spot-O-Matic has proven to be a more than capable tool for verifying that this and any future devices meet all of the SNAP pixel response requirements. The Spot-O-Matic has also proven to be a versatile piece of equipment, and has been able to perform a variety of additional detector measurements that will prove interesting and valuable to SNAP. These results are presented in the following sections.


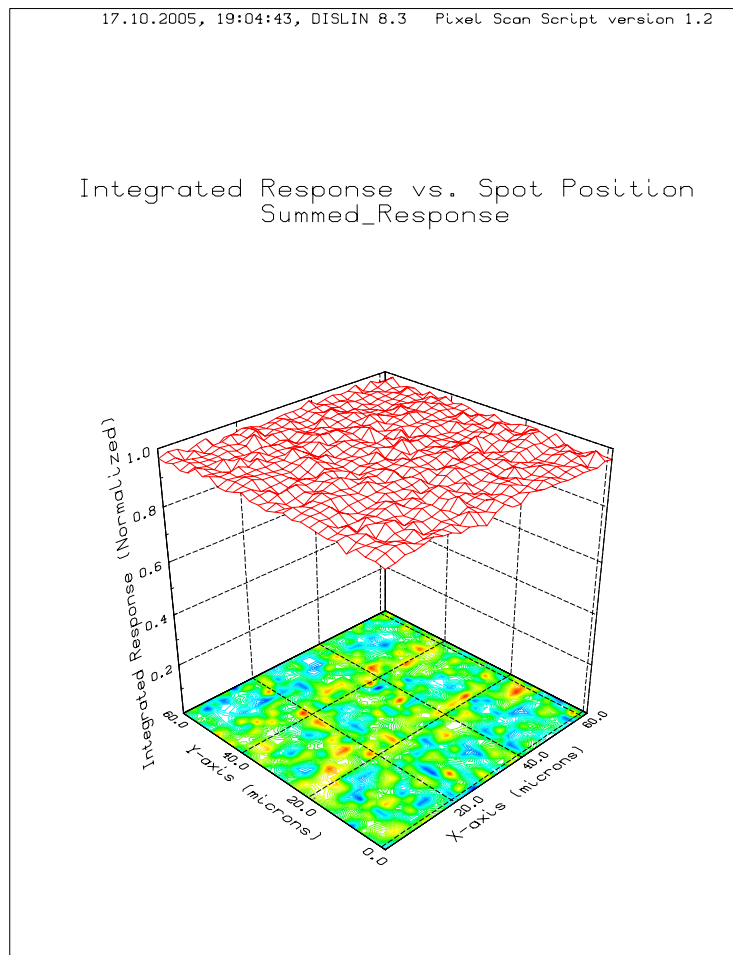


Figure 5.2: The top panel shows the sum of 9 adjacent PRFs scanned at 1050nm with the small diffraction-limited spot in Fig. 3.2. The response histogram is shown below.

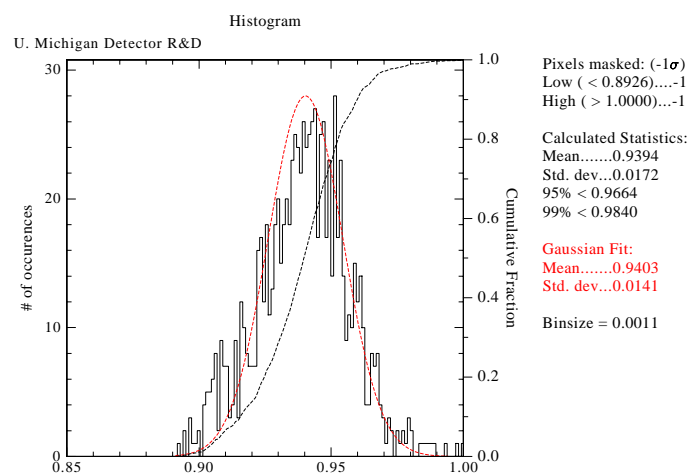



Figure 5.3: The top panel shows the sum of 9 adjacent PRFs scanned at 1550nm with the small diffraction-limited spot in Fig. 3.2. The response histogram is shown below.

5.4 SNAP Airy Disk Spot Scans

Until now it has not been discussed why it was acceptable to use the raw, convolved PRFs to determine the flatness of the device. Indeed, any features in the pixel response that were of similar scale or smaller than the spot used in the scan would not be resolved. Any fine structure would essentially be invisible to our scan. There could indeed be dips between pixels, but if they were no more than a micron wide, we would not see them. The reason that we are not concerned here is that SNAP will not see them either.

It was described in Section 1.2 that the SNAP optics will produce an Airy disk with FWHM of $10.2\mu m$ at $1000nm$. The scans described in Section 5.2 were performed with spots much smaller than this. Spots this small will be impossible within the SNAP optics. Since a scan done with a larger spot will resolve less fine structure and thus produce less variation, the plots in Fig. 5.3 and Fig. 5.4 show even more variation than will be resolvable by SNAP. To determine just how flat the device will appear to the SNAP optics, we will perform a scan with a spot of size similar to the SNAP Airy disk. It is quite possible that the bumps seen at $1550nm$ might in fact be invisible to the SNAP telescope. Since this Airy disk is the smallest possible structure that can strike the detectors, the real test of flatness is to ensure that a scan done with a spot *this* size has a variation under 2%.

In order to create a larger spot, we can either change the magnification by moving the pinhole closer to the tube lens, per Fig. 2.1, or use a larger pinhole. The housing of the tube lens prevents us from getting close enough to the center of focus to get a large enough spot, so we must instead use a larger pinhole.

A $100\mu m$ pinhole turned out not to be big enough, even at minimum magnification, and the only other size we had available was $1000\mu m$, so this is what we used. Since this pinhole is much larger, it must be placed further from the tube lens than the $10\mu m$ pinhole was, in order to get the correct image size. For our particular optics, trial and error was used to determine the correct distance from pinhole to tube lens for our desired spot size. The size of the spot is then measured using a standard real knife-edge scan. We verified that indeed the produced spots are very close to the size and shape of the Airy disk. The results are shown in Fig. 5.4.

We notice that the $1550nm$ projected spot fits much better to its respective calculated Airy disk than the $1050nm$ spot does. This is because we verified that the $1550nm$ spot fit the Airy disk first, and then simply switched the filter without moving the pinhole. As you can see, the calculated Airy disk gets much smaller when we switch to $1050nm$ light, but the projected spot does not. This is because the spot and pinhole are rather large, and therefore geometric optics dominate over diffraction. In order to make the $1050nm$ spot closer to the Airy disk, we would need to adjust the distance from the pinhole to the tube lens, as we did for $1550nm$. However, these spots are close enough for what we are attempting.

The scans are then conducted as before, with the exception that now with a $1000\mu m$ pinhole, we are dealing with 10000 times the flux as compared to the $10\mu m$ pinhole. We therefore place sufficient neutral density filters between the light-source and the liquid light-guide, in order to avoid saturating the device. The results of a scan done at $1550nm$ are shown in Fig. 5.5.

At first glance, the inter-pixel ridges do appear to be “smeared out,” as was expected. The response histogram in Fig. 5.5 shows a standard deviation of 0.0093, which is indeed

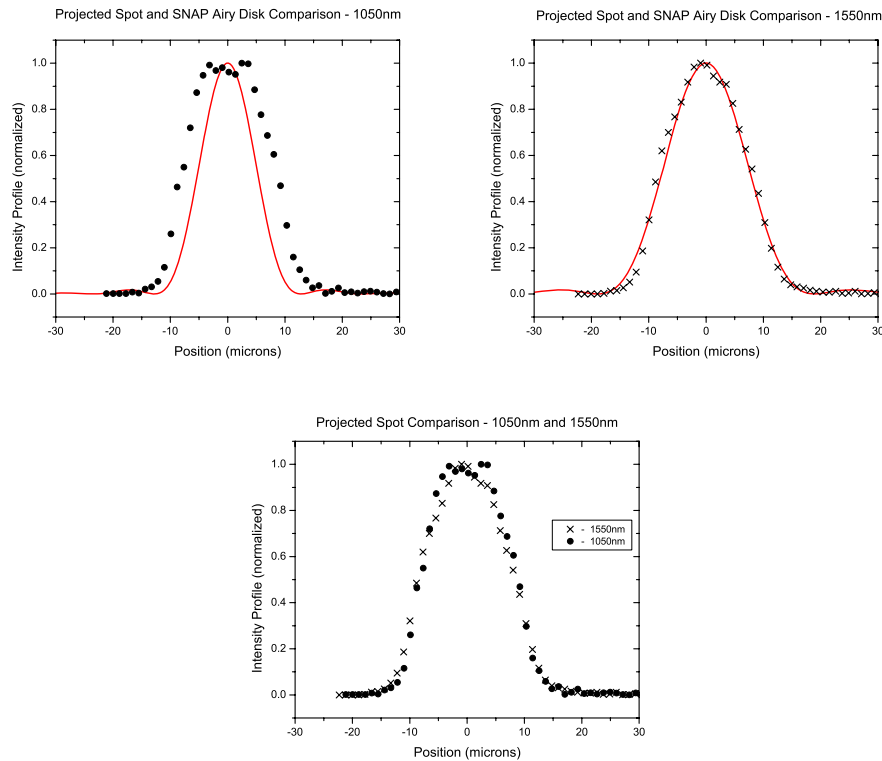


Figure 5.4: The top two panels show the knife-edge scan spot profiles created with the $1000\mu\text{m}$ pinhole at 1050nm and 1550nm . The red curve is the calculated Airy disk function for the SNAP telescope at the respective wavelengths. The bottom panel compares the shape of the measured spot profiles at each wavelength. This demonstrates that the discrepancy in fits in the top panels is due to the wavelength dependence of the calculated Airy disk, as described in Section 5.4.

significantly less than was observed when the pixels were scanned by the smaller spot at 1550nm . Unfortunately, due to time constraints on the InGaAs device, an additional scan at 1050nm could not be performed. However, even with the small spot the pixel response was virtually flat at this wavelength, so we should expect to see little change with a bigger spot. This scan at 1550nm does give us an upper limit on the pixel response variation that will be seen by the SNAP optics. This upper limit is quite acceptable, and it is possible that it could be lowered even further by correcting for non-linear response.

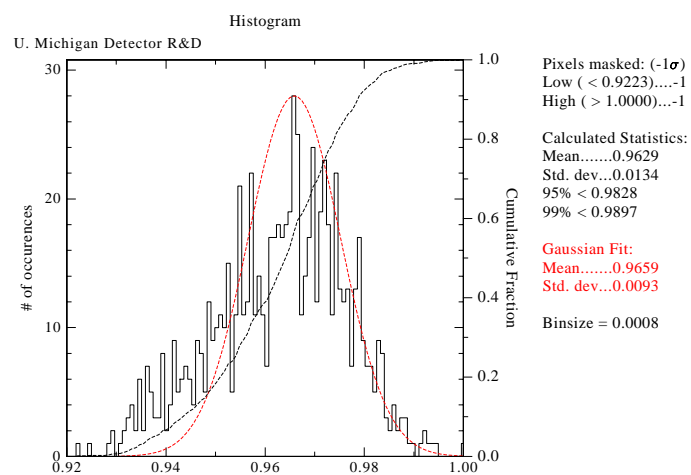
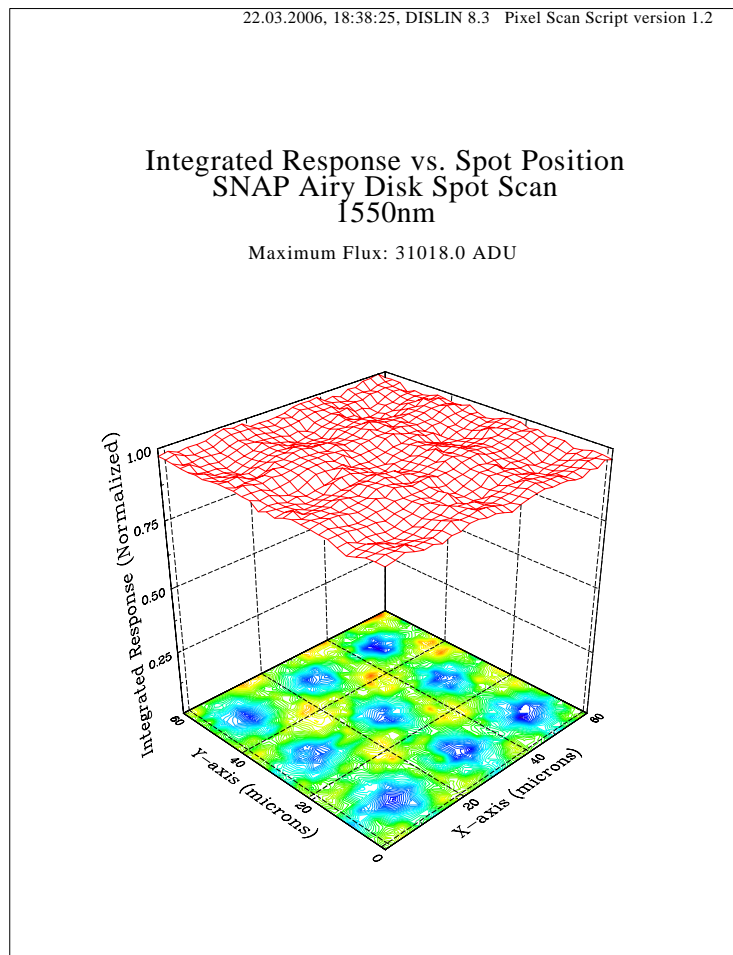


Figure 5.5: The top panel shows the sum of 9 adjacent PRFs scanned at 1550nm with the SNAP-like large spot in Fig. 5.4. The response histogram is shown below.

5.5 Dead Pixels

The InGaAs device, while performing well on the photometry tests, had a large number of surface defects (one of the reasons it was provided as a test device). Among these defects were a constant smattering of dead pixels. Under the exposure of background light, the dead pixels would exhibit no response. However, when we conducted a full scan over such a pixel, we discovered that it was not really dead. Fig. 5.6 shows a FITS image of the pixel of interest illuminated with background light. Fig. 5.7 shows the results of a scan of this single pixel. Fig. 5.8 shows the results of a summed scan of a region around the pixel.




Figure 5.6: FITS image of the dead pixel to be scanned.

The results are very curious for two reasons. First, it seems that we have stumbled upon two linked pixels. Fig. 5.7 shows the read-out of a *single* pixel vs. position, but it looks as though it were two adjacent PRFs added together. This means that these two pixels are sharing charge, not in the lateral charge diffusion sense, but rather their capacitors appear to be linked. The charge seems to be shared equally, which would suggest this conclusion. When we performed the same scan while reading-out *the other* pixel in the linked pair, we obtained identical results. It seems that the junctions of these two pixels became coupled in some way, essentially causing them to act as one big pixel.

The second and more puzzling result is that the pixel appears to be dead in Fig. 5.6, but we see that it is clearly not when it is scanned. Also, the base of the PRF for this pixel looks odd. It seems to abruptly cut off, instead of gradually tailing off as had the other pixels. This shape suggests that at a certain threshold of light intensity, this pixel suddenly “turns on” and begins functioning normally. This would explain why the pixel appeared dead when exposed to background light.

The summed scan in Fig. 5.8 shows that the overall QE is slightly lower for these linked pixels. It appears to be lower by approximately the amount missing from the base of the single pixel scan in Figure 5.7. The problem is believed to lie with the data acquisition electronics, and not with any physical defects in the detector. Such an effect could be observed if the bias voltage for this particular pixel (and others like it) was set incorrectly. The pixels in this device are set with an initial bias voltage of 0.5V. As the well fills up, an

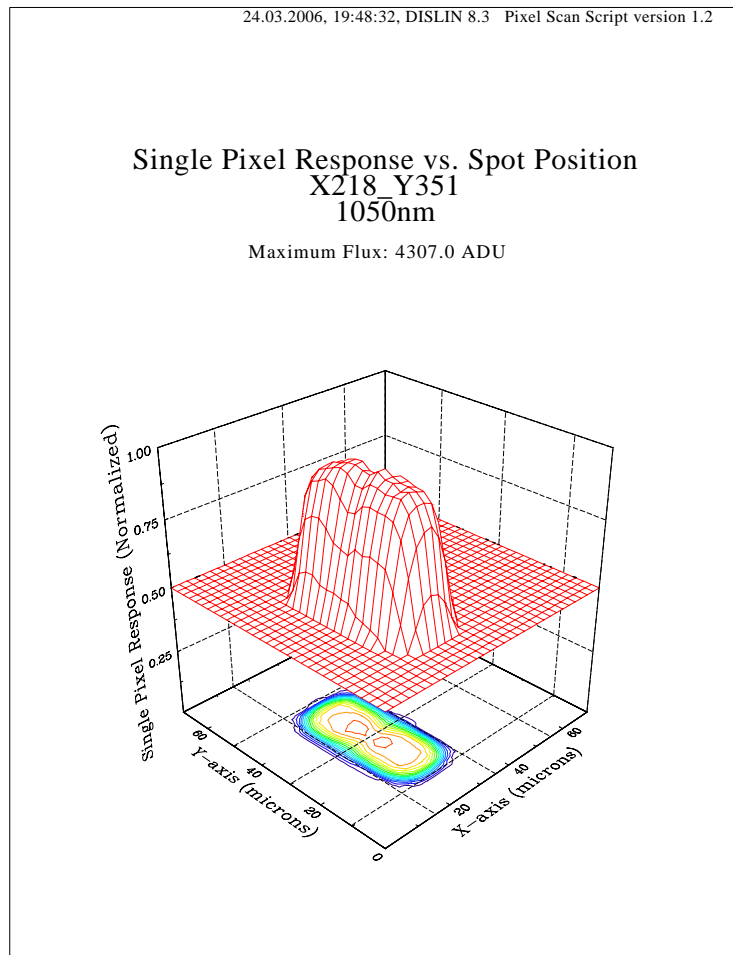


Figure 5.7: Response vs. spot position for an assumed dead pixel at 1550nm.

opposing electric field is created and this overall potential difference is read as light intensity in the pixel. The data acquisition electronics are set to read this initial bias voltage of 0.5V as zero intensity, and so if the bias voltage was initially set too high, the electronics would continue to read zero intensity until the well is sufficiently filled to bring the voltage back into the readable range.

This is not to say that all the dead pixels are not really dead. Some of these pixels truly are dead, and it is easy to see which they are. If a pixel is truly dead, the four adjacent pixels will read out slightly higher than the average pixel under a flat field. The simple reason is that any electrons in the conduction band of the dead pixel's substrate will have nowhere to go and will be collected into the neighboring pixels by lateral charge diffusion. Several of these dead pixels on a region of the device are shown in Fig. 5.9. Because of their appearance, we have named these features "daisies."




Figure 5.8: Summed pixel scan of the area around an assumed dead pixel at 1550nm.




Figure 5.9: FITS image of a region on the InGaAs with several “daisies,” illuminated with background light.

Chapter 6

Improvements

The Spot-O-Matic has proven to be an invaluable tool for very accurately measuring a wide variety of detector characteristics. There is still room for improvement however. The following sections detail improvements that have either already been made as of the writing of this thesis, or should be in the future.

6.1 Encoder Backlash

A linear encoder was installed on the z -axis to account for screw backlash, but the x and y axes still rely on the standard rotary encoders that were installed by the manufacturer. We have attempted to correct for this by always scanning from the same direction. This is not a perfect correction however, since it is clear from Fig. 2.4 that the backlash is not always the same amount, nor is the distribution necessarily symmetric. While this has not normally been a problem, it occasionally causes irregularities with some scans. Normally we simply throw these scans out, but if time constraints become an issue, this could present a problem.

Occasionally the drive screw on the stage becomes stuck momentarily, and then “jumps” back to where it should be several steps later. This failure to move is not picked up by the rotary encoders, and so we occasionally get scans that are “sliced” like the one in Fig. 6.1. It happens almost exclusively on the y -axis, since this is the axis that bears the Spot-O-Matic’s weight. If this becomes a problem in the future, it could be corrected by installing a linear encoder on the y -axis.

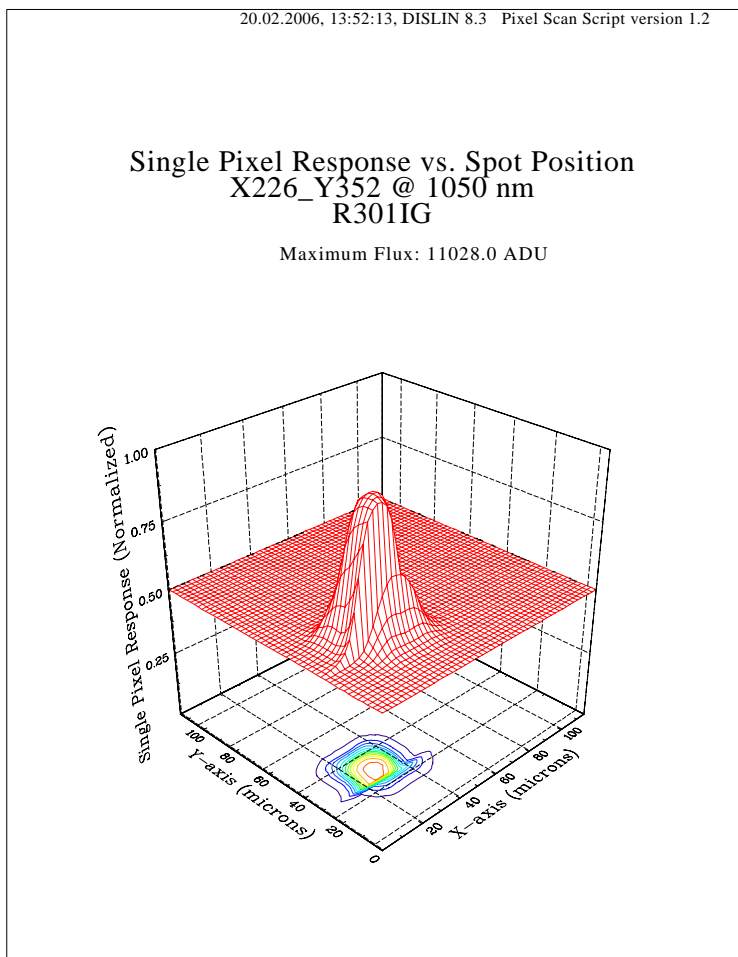


Figure 6.1: A single pixel scan with a measurement error caused by stage backlash.

6.2 Light Source Stability

As mentioned in Section 5.3, the scans shown in Fig. 5.2 and 5.4 are part of larger 5×5 scans that were not shown for clarity of experimental principle. The rest of the scans outside of the 3×3 region shown are not so flat or so pretty. These full scans and the corresponding response histograms are shown in Figures 6.2 and 6.3. The overall variation for this larger region is much higher, with standard deviations of 0.024 (Fig. 6.2) and 0.041 (Fig. 6.3) for the 1050nm and 1550nm scans, respectively.

The reason for the increased variation about the mean is clearly the response gradients across the scans. These gradients are troubling because they increase in the same direction that the scan progressed temporally, suggesting that they were not just a defect in the device. This must mean they are either due to variations in background, or variations in the light source intensity. It turns out it was a little of both.

Our dark box was not as dark as we had assumed, and it turns out that environmental variables like the lights being turned on/off in the lab or the sun rising/setting through the windows had a small but noticeable effect. Therefore, during future scans, the average response of the rest of the device not being scanned is recorded as a control, and then subtracted from each measurement. This was quite simple, and was not the main reason for the long-term intensity variation.

The problem lay with our light intensity controller. We could not find a manufacturer that made an intensity controller with an NIR diode, so we had to make due with the controller from Spectra Physics, which used a silicon diode that could not detect light above approximately 1100nm . We had placed a high bandpass filter with a 1000nm cutoff in front of the silicon diode, hoping that if we controlled the part of the spectrum closest to the NIR, we could keep the NIR spectrum from changing much.

The problem is that the 1000nm - 1100nm range lies on the front slope of the blackbody spectrum. As the lamp, the optics, and the tube assembly heated up over the course of a long scan, this would produce a significant increase in the 1000nm - 1700nm range that we are interested in, but only a modest increase in the range that we were monitoring. Therefore, the intensity controller would under-correct and the light intensity would be allowed to increase, causing the gradients that were seen.

Once this problem was realized, the intensity controller was turned off and future scans on devices after the InGaAs were performed with just the unregulated light source. Two of these scans are shown in Fig. 7.6 and 7.8. The light source has an ultra-stable power supply, and it seems that this was enough, as the scans came out quite flat, with no gradient.

However, it is clear from Fig. 2.2 that we are not making use of a potentially beneficial piece of equipment. If in the future we require more precision in our summed scans to more closely verify the SNAP photometry requirements, the intensity controller could be refitted with an NIR diode in order to function properly.

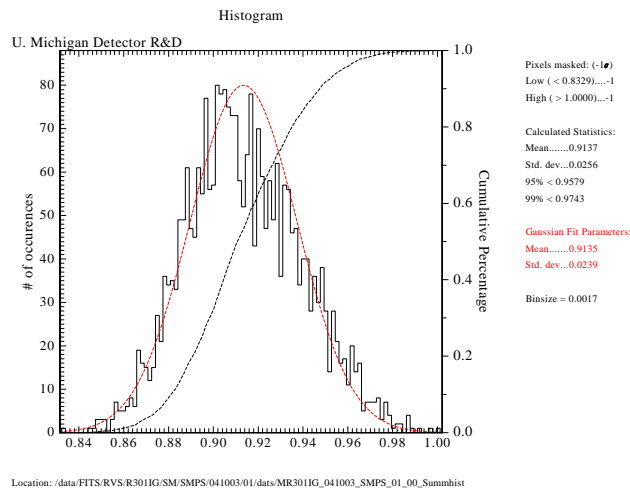
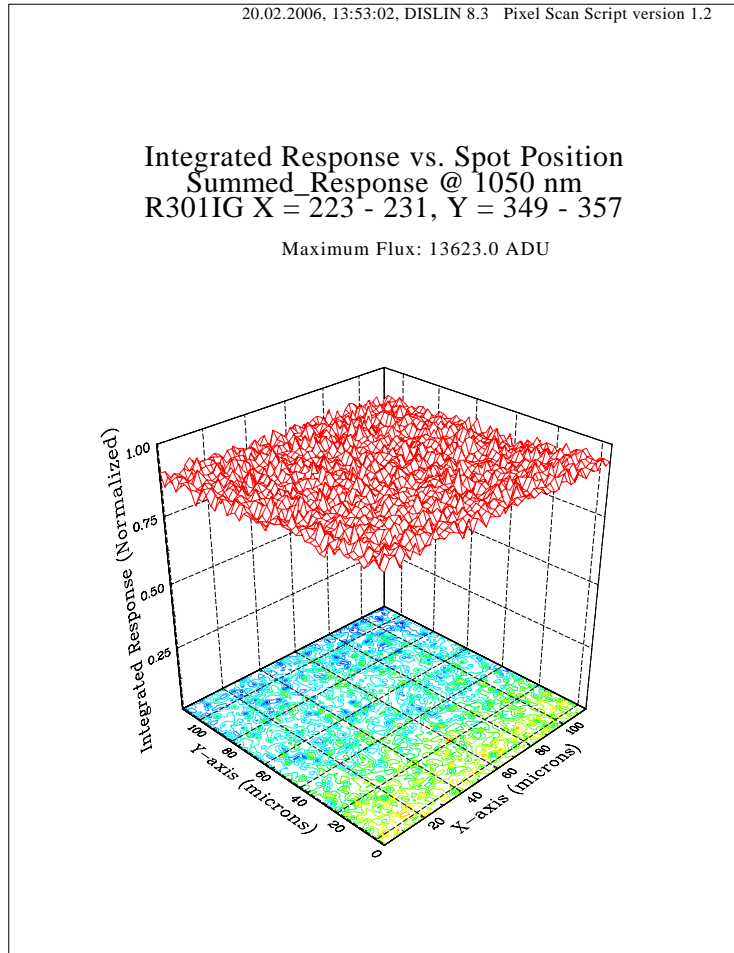


Figure 6.2: Full 5×5 summed pixel scan at 1050nm . Fig. 5.2 shows a subregion of this scan. The response histogram is shown below.

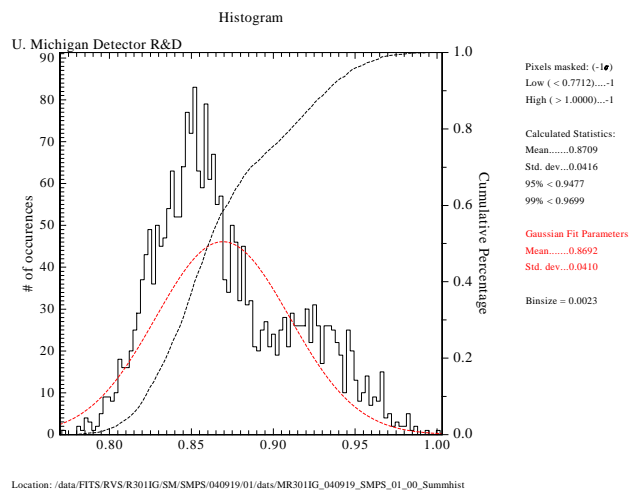



Figure 6.3: Full 5×5 summed pixel scan at 1550nm . Fig. 5.3 shows a subregion of this scan. The response histogram is shown below.

Chapter 7

Further Tests with the Spot-O-Matic

Since the completion of the tests on the Raytheon InGaAs device presented in this thesis, the Spot-O-Matic continues to be used by the University of Michigan SNAP group to characterize new devices. These devices are being tested with the ultimate goal of obtaining NIR detectors that are suitable to be sent up with the completed SNAP telescope. The Spot-O-Matic continues to produce reliable data with no degradation in precision or accuracy. It has produced invaluable and often unexpected insights into detector design and performance that has allowed the SNAP group to work with the manufacturers to produce better devices. The Spot-O-Matic continues to be the cutting-edge in NIR detector inspection. Some of the most recent and interesting results obtained by the SNAP group are summarized below. All data has been collected by Matt Brown of the University of Michigan.

7.1 Rockwell 1k banded HgCdTe Array

This test device from Rockwell was sent to the SNAP group to determine the results of changes in device construction. It did not perform well. The interpixel response was not nearly flat, and the individual pixel profiles were jagged and full of defects. Its only value is that it served as proof that the Spot-O-Matic could indeed detect such jagged structure on the sub-pixel level. A single pixel scan is shown in Fig. 7.1. A summed pixel scan is shown in Fig. 7.2.




Figure 7.1: Single pixel scan from the Rockwell 1k banded HgCdTe Array at $1550nm$.




Figure 7.2: Summed pixel scan from the Rockwell 1k banded HgCdTe Array at $1550nm$.

7.2 Rockwell H2RG HgCdTe #40

This device exhibited some very interesting and unexpected characteristics. The summed pixel scan came out with drastic peaks and valleys, although regularly and orderly spaced, as seen in Fig. 7.3.

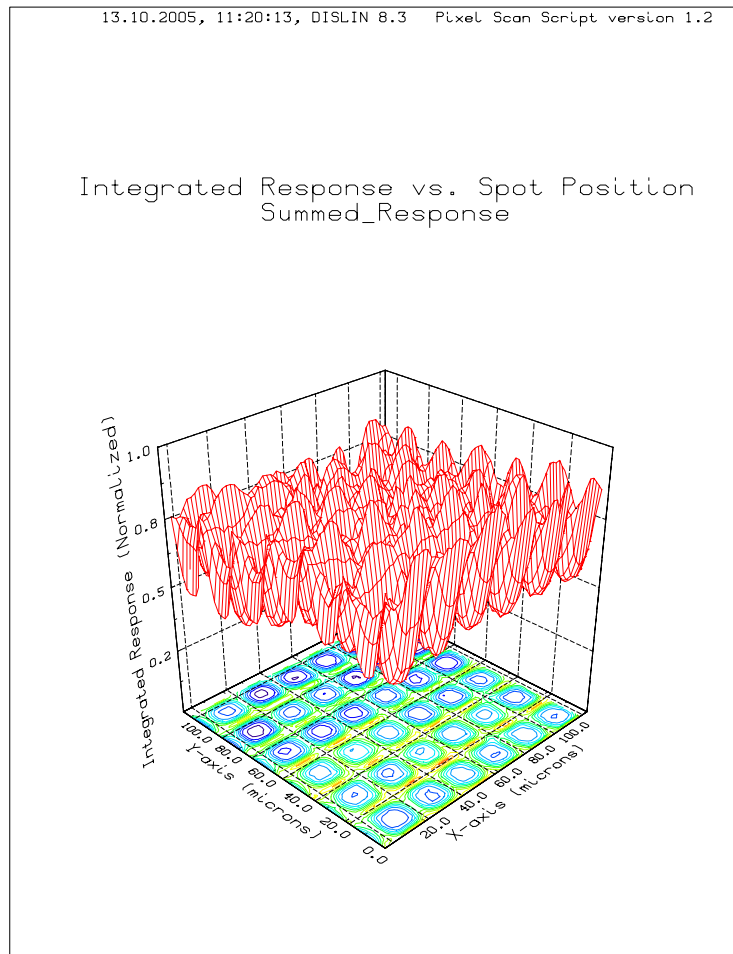


Figure 7.3: Summed pixel scan from the Rockwell H2RG HgCdTe #40 device.

When one looks at an individual pixel scan however, as in Fig. 7.4, the reason becomes clear. The pixel response of this device has one point of heightened sensitivity near one edge of the pixel, producing a “chair” shaped profile. It is in the same place for every pixel, leading us to suspect that this is where the pixel junction is located. If this is the case, this is exciting as it is the first time the Spot-O-Matic was able to “see” a junction. The virtual knife-edge scan in Fig. 7.5 shows this structure more clearly.

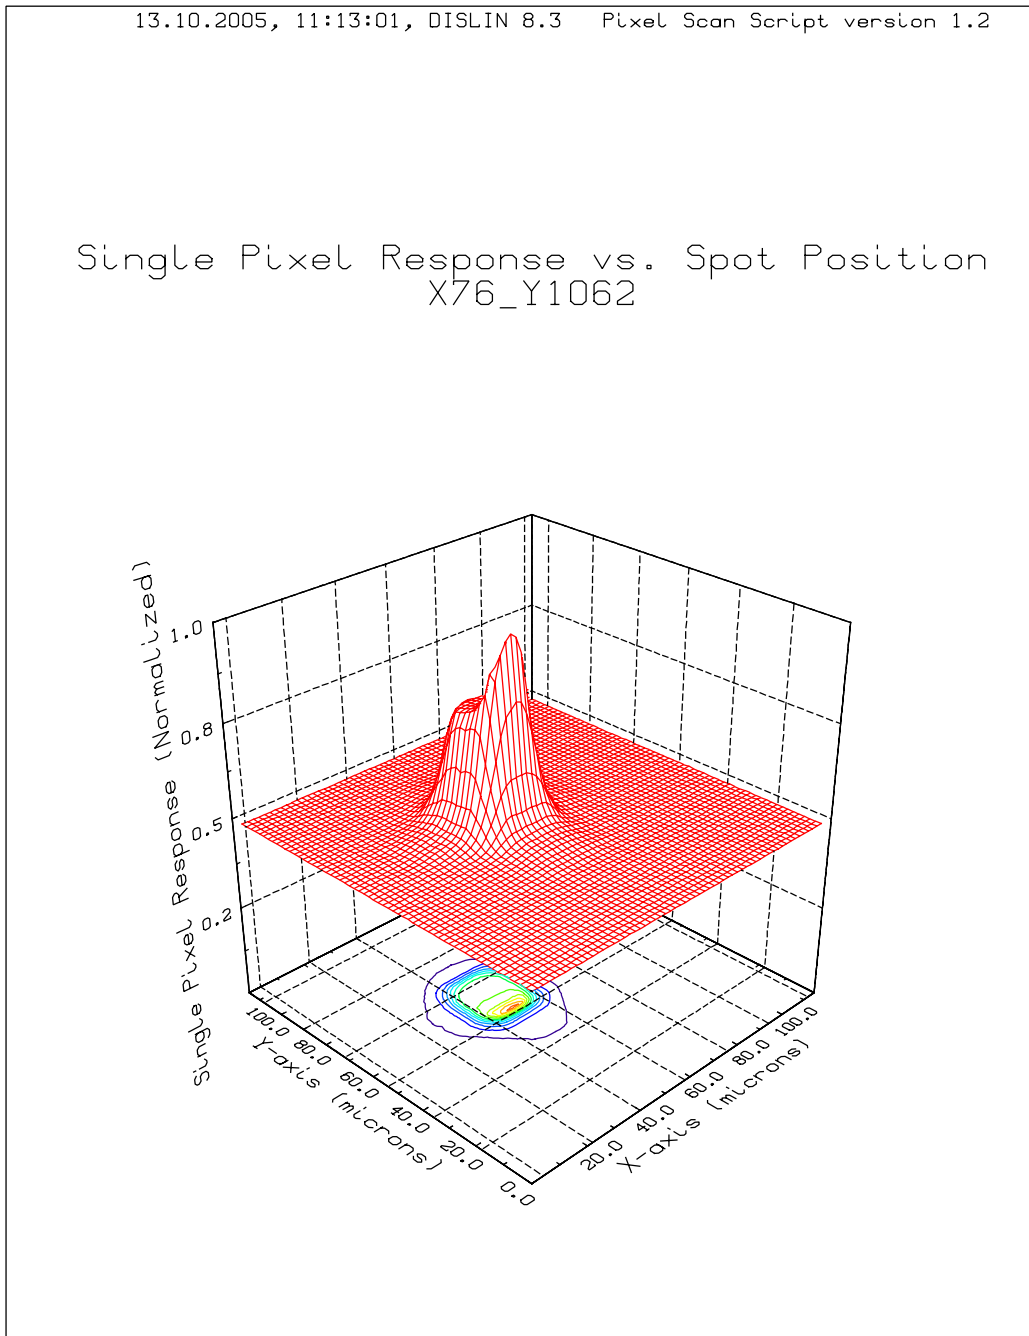


Figure 7.4: Single pixel scan from the Rockwell H2RG HgCdTe #40 device.




Figure 7.5: One-dimensional slice of a single pixel scan from the #40 device.

7.3 Rockwell H2RG HgCdTe #102

This device from Rockwell is notable for its good performance. Summed pixel scans and the corresponding response histograms are shown in Fig. 7.6 through 7.9. The summed pixel scans for 1050nm and 1550nm are both quite flat, with standard deviations of about 1.4% for each. These scans include a larger pixel region than the InGaAs scans (5×5 as opposed to 3×3) because the light intensity controller was turned off as described in Chapter 6, allowing us to perform longer scans without worrying about long-term light source variation. Thus a 1.4% variation over 25 pixels is excellent, and quite acceptable for SNAP. Several localized defects in the device are visible.

It is interesting to note that no difference in flatness is seen between the 1050nm and 1550nm scans. This is because this device does not have the same non-linear response problem as the InGaAs. However, the effect of lateral charge diffusion “smearing out” fine structure at lower wavelengths can be seen here as well. The defects are much more pronounced and sharp for the 1550nm scan, just as the inter-pixel ridges were on the InGaAs.

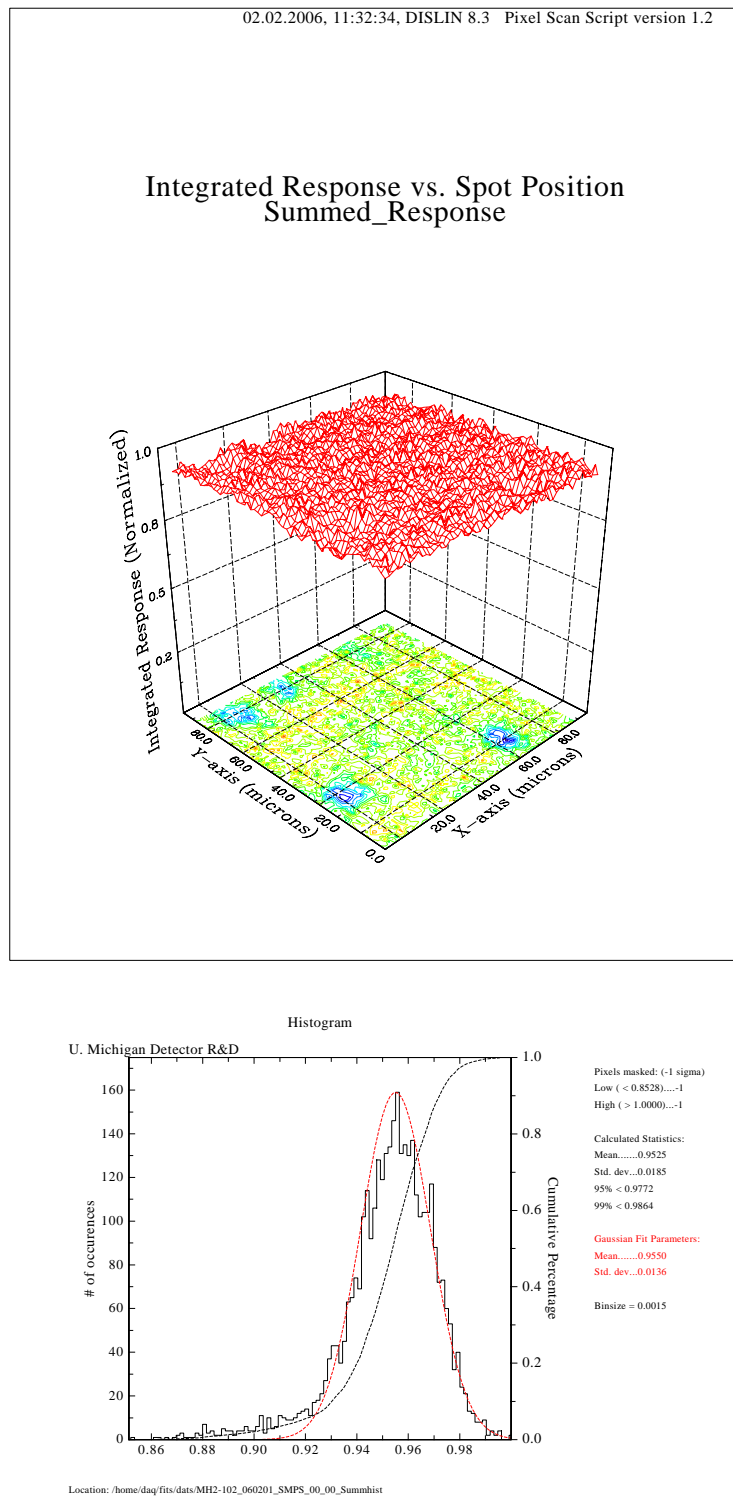


Figure 7.6: The top panel shows the summed pixel scan from the #102 device at 1050nm . The response histogram is shown below.


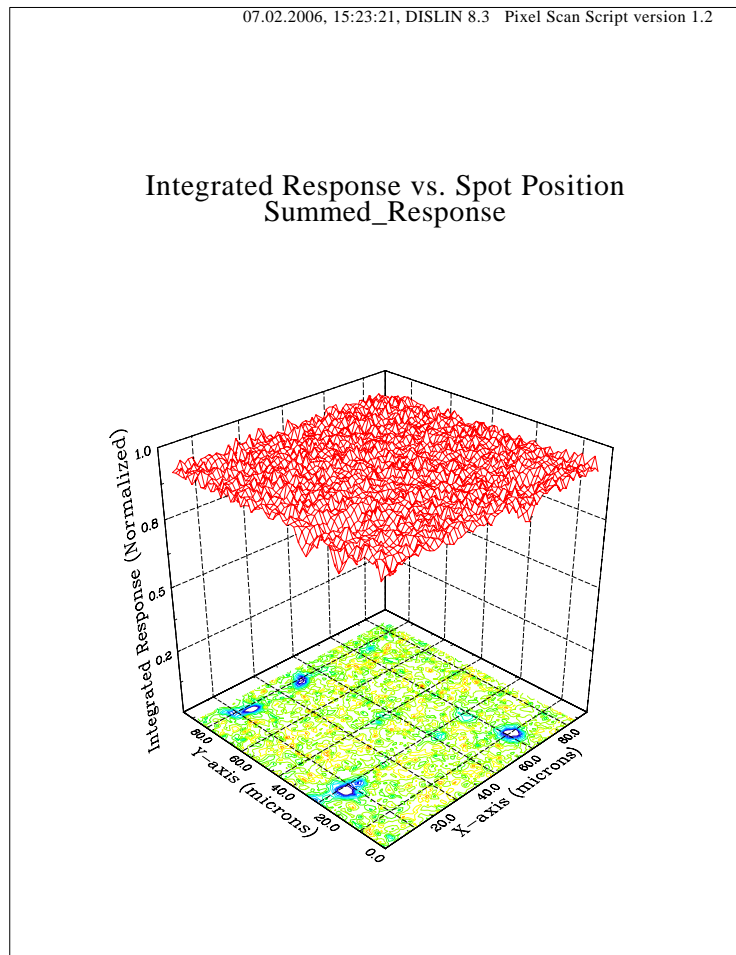


Figure 7.7: The top panel shows the summed pixel scan from the #102 device at 1550nm . The response histogram is shown below.

Bibliography

- [1] M. Lampton, et al, Proc. SPIE 4849 pp. 215-226, 2002.
- [2] G. Bernstein, PASP 114 p. 98, 2001.
- [3] Fruchter A.S. & Hook R.N., 'Drizzle: A Method for the Linear Reconstruction of Undersampled Images,' 2002, PASP, 114, 144.
- [4] M. Wagner, Thesis at the Lawrence Berkeley National Laboratory, 2002
- [5] Oriel Instruments, 'The Book of Photon Tools.'
- [6] Spectra-Physics, 300W Radiometric Power Supply user's manual, model #69931
- [7] Borysow M., 'Spot-o-matic: A Test Bed for Measuring Intrapixel Variation in Near Infrared Focal Plane Arrays.' 2004.
- [8] Goodliffe A.M., 'Time series analysis - theory and practice.'
<<http://www.geo.ua.edu/AMG/GEO369/class15.ppt>>
- [9] 'Suprnova/Acceleration Probe - Fact Sheet,' <<http://snap.lbl.gov/pdf files/snapfacts.pdf>>

Appendix A

Encoder Mounts

I designed the following two pieces in order to mount the linear encoder onto the Z axis of the stage. The encoder is the Mercury 3000 linear encoder from MircoE Systems with a $55mm$ long diffraction grating scale. The stage is the MM-4M-EX 3-axis stage from National Aperture. If machined to the specified tolerances, these mounts will perfectly align the encoder head with the diffraction grating for optimal performance. Be aware that these tolerances are extremely sensitive.




Figure A.1: Diffraction Grating Scale Mount

